

Essays on Small Business

Benjamin Wild Pugsley, Ph.D.

Abstract

Small businesses range from corner stores to high-tech startups and account for nearly all businesses in the United States. In part because of their small size, these businesses are intimately connected with their owners. This dissertation considers two important aspects of this connection. First, I study the role of the entrepreneurial tastes of small business owners in explaining the existence and variety among small businesses. Second, I consider the possibility of underreporting income both to tax authorities and to the surveys economists use to study entrepreneurship, savings, and income in the population.

Category: *Economics: Industrial Organization and Labor Economics*

Keywords: *Entrepreneurship, Firm Dynamics, Small Business Formation, Income Dynamics, Measurement Error*

Executive Summary

Introduction

Small businesses range from corner stores to high-tech startups. Firms with fewer than 20 employees account for nearly 90 percent of all firms in the U.S. and roughly 20 percent of total employment. Despite their relatively modest share of total employment, they represent the overwhelming majority of all businesses and encompass enormous variety. In part because of their small size, these businesses are intimately connected with their owners. This dissertation considers two important aspects of this connection. First, I study the role of the entrepreneurial tastes of small business owners in explaining the existence and variety among small businesses. Second, I consider the possibility of underreporting income both to tax authorities and to the surveys economists use to study entrepreneurship, savings, and income in the population.

The first chapter examines the economics underlying the formation, growth, and exits of businesses in the U.S. Modern theories of these dynamics link firm growth to productivity, access to credit, or some combination of the two. I present new evidence on persistent earnings differentials of business owners relative otherwise similar wage and salary workers that are difficult to explain with these two sources of heterogeneity. I extend the economic model of firm dynamics to incorporate varying preferences for business ownership consistent with the non-pecuniary motives documented in earlier work. Preferences turn out to be enormously important, explaining roughly 40 percent of small firms. I also show how ignoring preference heterogeneity leads to an overestimate, in particular for small businesses, of the tightness of borrowing constraints and a less realistic distribution of entrepreneurial ability.

The second chapter is joint work with Erik Hurst of the University of Chicago and Geng Li of the Federal Reserve Board. There is a large literature showing that the self-employed underreport their income to tax authorities. In this chapter, we quantify the extent to which the self-employed also systematically underreport their income in U.S. household surveys. To do so, we use a well documented relationship between income and expenditures (what economists refer to as the Engel curve) of wage and salary workers to infer the actual income, and thus the reporting gap, of the self-employed based on their reported expenditures. We find that, on average, the self-employed underreport their income by about 25 percent. This result is remarkably robust across data sources and alternative model specifications. We show that failing to account for such income underreporting leads to biased conclusions in existing measures of earnings and wealth differentials between the self-employed and wage and salary workers, precautionary savings, lifecycle earnings profiles, and earnings differences across the geographic regions.

Chapter 1 --- What makes small businesses small?

In this first chapter, I show that the standard economic models of firm dynamics (formation, growth or contraction, and exit) and the economy-wide distribution of firm size leave some key features unique to small businesses unexplained. These traditional accounts rely on differences in productivity or access to credit to explain firm dynamics and the observed firm size distribution. Of course, in practice there is tremendous heterogeneity along many dimensions among firms of all sizes and in particular in the small businesses that constitute the overwhelming majority of all firms. In any economic model, the goal is not to perfectly replicate all aspects of the real world. Rather, the goal is to identify and model the main driving forces of in this case firm dynamics, especially among smaller businesses.

I propose and evaluate a model with a new source of heterogeneity tied to the preferences of potential entrepreneurs. In earlier work with Erik Hurst at the University of Chicago published in the *Brookings Papers on Economic Activity* (Hurst and Pugsley, 2011) I document non-pecuniary reasons as the principal motivation among more than one-third of new entrepreneurs. These motives include a desire to be your own boss, pursue a creative career, and similar reasons unrelated to income, business opportunities, or a lack of employment options.¹ Those with non-pecuniary motives tended to have both less ambition for growth and less actual growth. I build from this insight and develop a model where the preferences of individuals are now relevant in the decision to become an entrepreneur and when to exit.

This chapter makes three scientific contributions. First, building on earlier empirical work, I document *persistent* earnings differences between otherwise similar business owners and workers. These differentials are particular to small business owners and are extremely persistent. Even adjusting for the underreporting, as I discuss in Chapter 2, the differentials create difficulties for the standard theories of firm dynamics. Second, I develop an alternative model that includes heterogeneity in the preferences of individuals towards running a business. I call this a “taste” for entrepreneurship. The model with tastes not surprisingly is better able to explain the earnings differentials. More importantly, incorporating tastes introduces a “selection on tastes” mechanism that turns out to be particularly relevant for understanding small firms. To better compare preference heterogeneity with alternative sources of heterogeneity, the model nests simplified versions of two popular models of firm dynamics that derive

from heterogeneity in productivity and in access to credit. Finally, using the model, I show that preference heterogeneity accounts for over 40 percent of small firms. I also show how ignoring tastes can overstate the role of credit constraints and productivity in understanding small firms.

Economics of the firm size distribution

Modern theories of firm dynamics build on two main sources of heterogeneity to explain the distribution of firm size. In the first, all firms produce the same good and firm size is purely a function of exogenous differences in efficiency.ⁱⁱ A model proposed by Hugo Hopenhayn (1992) of firm dynamics and its derivatives are prime examples. In these models, firms often have decreasing returns to scale and a fixed cost. Productivity fluctuates over time and firms exit the market when they are no longer profitable. The decreasing returns or similar feature ensure an optimal scale, because costs grow too rapidly beyond a certain point that depends on firm productivity (or efficiency). The fixed cost delivers a minimum firm size because profits below a certain size are too small to offset the fixed cost. This fixed cost could be technological as in the original Hopenhayn (1992) model, or it could follow from an occupational choice as it does in this paper; running the firm requires time that could otherwise be allocated to the labor market. In efficiency based models, the distribution of firm size follows entirely from the underlying distribution of productivity. Larger scale firms have lower costs than small firms.

In the second kind of model, differences in firm size follow from the wealth distribution of households. In these models, financial market imperfections, such as imperfect enforceability of contracts, require the firm post collateral in order to borrow capital. The credit constraints prevent firms from borrowing enough capital to immediately reach their efficient scales. Wealthier entrepreneurs are able to run larger firms because they have the resources to borrow larger amounts of capital. In Cagetti and De Nardi (2006), for example, all entrepreneurs are equally efficient, and the distribution of firm size follows entirely from the distribution of wealth.

These types of models, often in combination, have had enormous success in explaining firm entry, exit and growth, as well as explaining the savings patterns of business owners. Cooley and Quadrini (2001) combine both sources of heterogeneity to explain age and size dependence in firm growth. Moll (2010), Buera et al. (2011), and Midrigan and Xu (2010) have all used the combination of heterogeneity in productivity and credit availability to measure the extent of capital misallocation when highly productive entrepreneurs cannot borrow or are deterred from starting a business. Collectively, the implication of these models is that firms are either small because they are less productive or because they are more productive but slowly growing into their efficient scale.

Reconciling firm dynamics with the earnings evidence

When posed with an occupational choice decision between employment and entrepreneurship, these models imply selection solely on comparative advantage. Entrepreneurs choose to run businesses because of a higher expected rate of return on their labor and financial capital than in employment and the savings market. This is harder to support empirically. I present evidence that business earnings are in fact almost 30 percent lower than comparable worker earnings for the typical business owner. Although, previous work by Hamilton (2000) and Moskowitz and Vissing-Jørgensen (2002) has documented low returns from business ownership overall, I examine the heterogeneity underlying the aggregate returns .

I show these earnings differences are driven largely by owners of small firms. Moreover, they persist throughout the firms' life cycle, even as uncertainty over firm performance fades. Even after adjusting for potential underreporting, heterogeneity in productivity or access to credit cannot explain these facts. When employment is a viable outside option, entrepreneurs need sufficient compensation to forgo employment and form a business. In some examples of these models, given the risks of business ownership, this additional compensation may be substantial.

To confront these earnings differentials I develop a model that introduces preference heterogeneity towards running a business. The model extends work by Midrigan and Xu (2010) and constructs a Hopenhayn and Rogerson (1993) style general equilibrium that nests versions of efficiency- and credit-heterogeneity driven firm dynamics. People vary in their tastes for running a business, which are fixed forever. For those with a taste for entrepreneurship, all things equal, they would prefer entrepreneurship to employment.

Given a set of skills (ability or productivity as an entrepreneur and productivity as a worker) that evolve stochastically, people make an occupational choice each period to either run a business or work for a business run by another household. All businesses produce the same good with a decreasing returns to scale technology. This construction, similar to the Hopenhayn (1992) style models, ensures an optimal size for the business and a threshold where it is no longer desirable to operate the business. This threshold depends on the value of switching back to employment and the tastes, since returning to employment means forgoing the additional non-pecuniary compensation of entrepreneurship. In the absence of credit frictions, solely the ability of the business owner would determine the size of the firm. However, I include an exogenous collateral constraint that limits the debt the firm can take on and by extension limits the size of the firm. Credit constrained firms must grow into their optimal scale by accumulating additional savings as collateral.

The key contribution of the model is to provide a role for preferences in the occupational choice decision. Business owners are driven by comparative advantage and tastes. Incorporating preferences changes the composition of operating businesses by lowering both the minimum wealth and ability levels needed to run a business. As a result, a set of businesses operate that would otherwise not have formed. All of the marginal firms are small scale businesses, which are small because their size is already the optimal scale. These businesses exist alongside higher ability but credit constrained firms growing into their optimal scale, resembling the heterogeneity among small businesses in the United States.

The marginal business owners will have lower pecuniary earnings than comparable workers in return for the non-pecuniary compensation generated by their preferences. Since marginal businesses are smaller scale, they are less likely to be credit constrained. Whereas credit constrained firms have high expected returns to match their higher borrowing costs, the expected return of marginal firms is close to the risk free rate.

Non pecuniary motives are a powerful force

In a calibrated version of the model that matches the patterns of earnings differentials, I find that taste-driven marginal firms account for over 40 percent of all firms. Although previous work has

acknowledged the possibility of non pecuniary compensation of business owners, this paper is the first of its kind to quantify the role of tastes in the composition of small businesses. By letting tastes interact with the standard forms of heterogeneity, which we know to be extremely important, I show the consequences of ignoring tastes. This addresses an important challenge with integrating an occupational choice into models of firm dynamics.

Without selection on tastes, we need another way to entice proprietors of small firms to stay in business in spite of more lucrative employment earnings. Purely credit constraint explanations require these small firms to be constrained with high expected returns. To achieve this with only productivity requires a stochastic process that makes earnings only temporarily low. However, in both cases, low earning small scale firms cannot survive in the long run when there is a viable outside option.

The change to the selection mechanism depends in part on tastes that do not scale with the size of the business. Preferences shift the occupational choice decision without affecting the investment and employment decisions as a business owner. This comes at some cost complicating the solution of the model. However, making tastes independent of business size accounts well for the earnings differentials of small business owners.

As an additional evaluation of the model and this form of preference heterogeneity, I consider an extension to two sectors with different levels of fixed costs so that each sector has a different average scale. Suppose business owners are equally capable of running firms in either sector given their ability. The implication of selection on size-independent tastes is that the sector with lower fixed costs will contain a higher proportion of taste-driven business owners, and thus the earnings differentials relative to working will be larger. I find this same pattern empirically. Earnings differentials in the manufacturing sector with larger fixed costs are much smaller than the service sector with smaller fixed costs.

Finally, the model helps compare the role of differences in entrepreneurial tastes with differences in productivity or access to credit. As discussed, models without non-pecuniary compensation have a difficult time explaining small businesses in the long run. In these models, small businesses are only temporarily small. They may expect to be more productive, or even believe they are more productive than they actually are. In these cases, the business will grow or exit. Alternatively, they may be highly productive but without sufficient savings to use as collateral, in which case the business will save its way into a larger more profitable firm. With the taste differences eliminated, to create enough small businesses, credit constraints must be an order of magnitude more severe. Even in a model with no collateral constraints, the distribution of entrepreneurial ability must be less dispersed. A more realistic ability distribution with a long tail would leave only a small number of very high ability entrepreneurs employing everyone else. Introducing entrepreneurial tastes restores the incredible diversity among actual small businesses in the U.S. with some young small firms growing quickly towards their high potential and generating high returns and some persistently smaller businesses unconstrained and content to remain small.

Chapter 2 --- Difficulties of Measuring Small Business Earnings

The second chapter of my dissertation, written in collaboration with Erik Hurst at the University of Chicago and Geng Li at the Federal Reserve Board studies an additional characteristic of small businesses in the U.S., their owners' ability to underreport income to the IRS. There is already a large body of empirical evidence confirming individuals' underreporting income to the tax authorities.ⁱⁱⁱ Even other sources of administrative data, such as the Vital Statistics, are susceptible to misreporting. For example, Black et al. (2009) compares administrative data from Vital Statistic records to data from the U.S. Census to show that in states where minimum age of marriage laws were binding, younger individuals appeared to have lied about their age to government officials when applying for their marriage license. A completely distinct experimental economics literature finds that participants in experiments distort their behavior as a reaction to being studied.^{iv} Collectively, previous research has shown that individuals tend to misreport their actual behavior to data collectors and administrative agencies when the incentives to do so are sufficiently large (e.g., to avoid tax payments) and/or the cost of doing so is small (e.g., changing their behavior in experimental settings).

An implicit assumption made in the majority of empirical work using household survey data, however, is that the data within the household surveys are immune to such systematic errors. In doing so, researchers are assuming that the problems that plague tax data, experimental data, and other types of administrative data do not plague household survey data. In this chapter, we assess whether such an assumption is valid. Specifically, we address whether the self-employed, who have been shown to have misreported their income to tax authorities, have also misreported their income to household surveys.

A natural question that arises is why would someone misreport their earnings to household surveys? On the one hand, there is very little direct benefit for an individual to underreport income to household surveys. Unlike when reporting to tax authorities, misreporting income to household surveys does not decrease the individual's tax burden. On the other hand, there is also no penalty associated with misreporting income to household surveys.

Individuals, particularly the self-employed, may perceive other indirect net benefits from underreporting their income to household surveys. First, unlike wage and salary employees who receive W-2 forms, the self-employed have to expend effort to accurately account for their true income. If the self-employed have already supplied (or are intending to supply) a distorted income report to tax authorities, it may be easier to use this distorted report when responding to household surveys instead of taking the time to compute a second and more accurate measure of income. In fact, many household surveys recommend that households consult their tax records when completing the survey. Second, because some judgment may be involved in the computing of self-employed income, individuals may believe that the income report provided to tax authorities which reduces their tax burden is actually their true measure of income. Finally, households may feel compelled to provide consistent measures of their income between the tax authority and household surveys if they believe that their reports to household surveys may not be completely confidential. Given that the self-employed already have strong incentives to misreport their income to tax authorities and given that they face no penalty for misreporting to household surveys, even a small probability of self-incrimination may be enough to distort their survey responses. An extreme form of the last story is not likely driving our results. Individuals who feel

compelled to provide consistent measures of their income to household surveys to minimize self-incrimination would also have incentives to underreport their expenditures to household surveys given the close link between income and expenditure. As we show, we find no evidence of systematic expenditure underreporting by the self employed

As a result of any of the above three stories, underreporting of income to tax authorities may simply be passed along when the individual reports their income to household surveys. With no penalty compelling an accurate response but potential net benefits from providing a distorted one, economic theory predicts that households, particularly those who already underreport their income to a tax authority, may also underreport their income to household surveys. The extent of underreporting of income to household surveys, therefore, becomes an empirical question.

An expenditure-based estimate of earnings

We infer using expenditures the extent to which the self-employed underreport their income within U.S. household surveys. As far as we know, this is the first paper that attempts to do so. To estimate the extent of underreporting, we follow a procedure similar in spirit to the one set forth in Pissarides and Weber (1989). This procedure estimates the relationship between expenditures and income for wage and salary workers and uses the estimated coefficients from this relationship to predict the true income of the self-employed based on their reported level of expenditures. To estimate the relationship for wage and salary employees, we must account for the fact that expenditures are determined both by current income and expectations of future income. Once properly estimated, we ask in essence, how much would a wage and salary employee need to earn on average to generate the level of expenditures we see among a similar self-employed individual?

We use data from two well-established and long-running household surveys: the Consumer Expenditure Survey administered by the U.S. Bureau of Labor Statistics and the Panel Study of Income Dynamics administered by the University of Michigan. We find that, on average, the self-employed underreport their income to household surveys by about 25 percent. The estimated magnitudes are nearly identical across both surveys for similarly defined specifications.

Our estimating procedure makes a few key assumptions. First, it assumes no differential underreporting of expenditures by the self-employed relative to wage and salary workers. If the self-employed underreport income, they might also adjust their reported expenditures to maintain parity. This would make us underestimate the extent of underreporting. Second, it assumes that the conceptual measures of income in household surveys are similar between wage and salary workers and the self-employed. Finally, it assumes that the underlying relationship between income and expenditures, absent any underreporting, is similar between wage and salary workers and the self-employed. If the self-employed spend more, perhaps because some personal consumption expenditures might be expensed to the business, we would overstate the extent of underreporting. We test for the validity of all of these assumptions, as well as provide a variety of additional robustness analyses.

Accounting for underreporting matters

A second goal of the chapter is to show several examples of how ignoring the underreporting of income by the self-employed can bias many different types of empirical analyses. Not surprisingly, we show that studies that focus on earnings or savings differences between the self-employed and wage and salary workers are severely biased if the underreporting of income by the self-employed is not accounted for. For example, Hamilton (2000) finds that the earnings of the self-employed are 35 percent lower than that of otherwise similar wage and salary workers. However, accounting for the underreporting of income by the self-employed would dramatically reduce these estimated earnings gaps. Adjusting for the misreporting of income by the self-employed also reduces the well documented estimated wealth gaps (conditional on earnings and demographics) between the self-employed and wage and salary workers by roughly one third.

We also show that not accounting for the underreporting of income by the self-employed can lead to quantitatively different results in a variety of other settings. For example, we show that estimated results on the importance of precautionary savings are 10 percent lower when we take into account the extent to which the self-employed underreport their income. Similarly, given that self-employment propensities differ over the lifecycle and across space, measures of lifecycle and spatial differences in earnings are sensitive to the extent of underreporting of income by the self-employed. For example, we find that roughly between 5 and 10 percent of the decline in earnings between the ages of 45 and 65 in both the Consumer Expenditure Survey and the Panel Study of Income Dynamics can be explained by the underreporting of income by the self-employed given that self-employment propensities rise with age.

On balance, our work shows that it is naive for researchers to assume that individuals willing to distort the truth in some contexts would automatically provide accurate responses when participating in household surveys. While the benefits of providing distorted information to household surveys are small, so are the costs of providing inaccurate information. Such potential biases need to be accounted for when analyzing data from household surveys.

ⁱ The full list of non-pecuniary reasons is: be own boss; tired of working for others; Flexibility; more free time; set own hours; Stay home with children; work from home; Enjoy the work, have passion for it ; hobby; Job Security/Financial independence; Try new career; charge career; do something new; creative ; do creative work; creative outlet; better life; lifelong ambition; challenge; personal growth; to do more fulfilling work; other lifestyle references; other work desirability references

ⁱⁱ Although I focus on productivity and credit frictions, there are many other potential sources of differences among firms. Rossi-Hansberg and Wright (2007) consider differences in the accumulation of an industry specific form of human capital as a factor of production. See Luttmer (2010) for an excellent summary of models of firm dynamics.

ⁱⁱⁱ See, for example, Clotfelter (1983); Slemrod (1985); Feinstein (1991); Andreoni et al. (1998); Slemrod (2007); Feldman and Slemrod (2007).

^{iv} This phenomenon is sometimes referred to as the Hawthorne Effect, named for a series of studies at the Hawthorne Works factory where workers' productivity initially improved while under observation but declined soon after. There is a large literature within the social sciences on the Hawthorne Effect. See Levitt and List (2011) for a recent discussion.

References Cited

- Black, R., Charles, K., & Saltee, J. (2009). A Cautionary Tale about the Use of Administrative Data: Evidence from Age of Marriage Laws. *American Economic Journal: Applied Economics* , 1 (2), 128-49.
- Buera, F. J., Kaboski, J. P., & Shin, Y. (2011). Finance and Development: A Tale of Two Sectors. *American Economic Review* , 101 (5), 1964-2002.
- Cagetti, M., & De Nardi, M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy* , 114 (5), 835-870.
- Cooley, T., & Quadrini, V. (2001). Financial markets and firm dynamics. *American Economic Review* , 91 (5), 1286-1310.
- Hamilton, B. (2000). Does entrepreneurship pay? An empirical analysis of the returns to self-employment. *Journal of Political Economy* , 108 (3), 604-631.
- Hopenhayn, H. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society* , 60 (5), 1127-1150.
- Hopenhayn, H., & Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy* , 101 (5), 915-938.
- Hurst, E., & Pugsley, B. (2011). What do Small Businesses Do? *Brookings Papers on Economic Activity* , Fall.
- Midrigan, V., & Xu, D. (2010). *Finance and misallocation: Evidence from plant-level data*. Tech. rep., National Bureau of Economic Research.
- Moll, B. (2010). *Productivity Losses from Financial Frictions: Can Self-financing Undo Capital Misallocation?* unpublished.
- Moskowitz, T., & Vissing-Jorgensen, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review* , 92 (4), 745-778.
- Pissarides, C. A., & Weber, G. (1989). An expenditure-based estimate of Britain's black economy. *Journal of Public Economics* , 39 (1), 17-32.