



Dissertation Advisors: Lawrence Katz  
Edward Glaeser, Josh Lerner

Naomi Hausman

## Essays in Entrepreneurship and Public Economics

### ABSTRACT

This dissertation consists of three essays in entrepreneurship and public economics.

Universities are believed to be important drivers of local economic growth. The first essay identifies the extent to which U.S. universities stimulate nearby economic activity using the interaction of a national shock to the spread of innovation from universities – the Bayh-Dole Act of 1980 – with pre-determined variation both within a university in academic strengths and across universities in federal research funding. I find that long-run employment and payroll per worker around universities rise particularly rapidly after Bayh-Dole in industries more closely related to local university innovative strengths. The impact of university innovation increases with geographic proximity to the university. Entering establishments - in particular multi-unit firm expansions - over the period from 1977 to 1997 were especially important in generating long-run employment growth, while incumbents experienced modest declines, consistent with creative destruction. Suggestive of their complementarities with universities, large establishments contributed more substantially to the total 20 year growth effect than did small establishments.

The second essay measures capital gains tax lock-in for household portfolios and estimates the deadweight cost of this behavioral response. Taxes inhibit about one-fourth of the rebalancing in which households would otherwise engage, resulting in a simulated excess burden of 10-25% of revenue for the median household.

The third essay evaluates a common argument in the US health policy debate that rising health insurance costs, coupled with the tying of insurance to employers, inhibits the survival and growth of entrepreneurial firms. Economic theory suggests that these costs may indeed adversely affect small businesses: the firms may be unable to pass on to employees the full cost of benefits due to downward nominal wage rigidities or labor market competition with large firms. Instrumenting for premium growth in multiple ways, I can rule out large negative effects of rising health costs on firm survival and employment growth. I also find that firms facing higher growth in health costs are more likely to offer insurance, possibly due to the higher value of insurance when potential loss rises.

Essays in Entrepreneurship and Public Economics  
EXECUTIVE SUMMARY  
Naomi Hausman  
June 2011

## **Effects of University Innovation on Local Economic Growth and Entrepreneurship**

What is the effect of universities on neighboring industry? Many of the most innovative and entrepreneurial places in the United States cluster around research universities, such as Silicon Valley around Stanford and Boston's Route 128 corridor around Harvard Medical School and MIT. Such agglomeration of industrial activity may arise for a number of reasons, including shared inputs, local natural advantages, skilled labor pooling, and, importantly, knowledge spillovers (Marshall (1890); Krugman (1991)). Proximity facilitates increased interaction between people and speedy flow of ideas, compounding the positive externalities of knowledge production and encouraging locally concentrated growth (Jacobs (1969); Lucas (1988)). Universities and research hospitals – important generators of new ideas and disposed towards openness in discovery – are thus natural suspects as contributors to the local economy (Jaffe (1989); Furman and MacGarvie (2007)). But feedback effects from business activity and common underlying factors affecting both universities and industry make universities' influence difficult to measure. This paper uses a new strategy based on universities' technological strengths and a change in federal policy to identify effects of universities on the growth of neighboring industry.

Before 1980, universities lacked strong incentives to commercialize research; the federal government held rights to all intellectual property produced by universities in the course of federally funded research.<sup>1</sup> But the December 1980 passage of the Bayh-Dole Act gave

---

<sup>1</sup> Of course, some universities, especially public and land grant, were more practically oriented long before Bayh-Dole (Mowery et. al. 2004; Sampat 2006; Goldin and Katz 2008). Because they were so heavily funded by the

universities property rights to innovations developed under federal funding, and with these strong new incentives it opened a sea of patenting and licensing activity from universities as they developed infrastructure for technology transfer that they previously lacked (Henderson, Jaffe, and Trajtenberg (1998); AUTM Licensing Activity Survey).<sup>2</sup> That this legal change fundamentally increased universities' connection to industry and induced greater spread of innovation from universities will help to identify their local effects.<sup>3</sup>

My strategy interacts the law change with cross-sectional variation in the extent to which industries benefit from nearby university innovation. Because universities produce more innovation – as measured by patents – in some technological areas than in others, I am able to identify the industries surrounding each university that are most likely to benefit from this increased spread of innovation. The industry intensities provide variation both within university, between industries, and between universities due to their different industrial mixes. This type of variation has the advantage that I can hold a geographical location fixed and identify an effect off of cross-industry differences in the intensity of field-specific innovation from the nearby university. I am further able to address the concern that universities may simply innovate in the nation's most quickly growing industries, such as biotech in the late 1980s and early 1990s, by controlling for nationwide changes in industry performance.

---

states, they always faced incentives to be responsible to local industry. Even these institutions, however, tended to keep their commercial arms divorced from the university, administering patents through research foundations (like Wisconsin Alumni Research Foundation) or third parties such as the Research Corporation.

<sup>2</sup> New Congressional endorsement of the value of these activities to the economy may have also been important in changing the anti-commercialization sentiment harbored in many universities.

<sup>3</sup> There is some dispute over the degree to which Bayh-Dole altered research and commercialization practices in universities. Mowery and Ziedonis (2000) and Mowery et. al. (2001) find in case studies of three universities that Bayh-Dole had little impact on the content of research, substantial impact on marketing efforts, and little impact on final patenting and licensing outcomes. More comprehensive patent data indicate, to the contrary, a substantial increase in university patenting after Bayh-Dole, though this increase may reflect higher production of innovation, increased attention to commercialization of innovation, or both. By showing differential industry growth around universities after the law relative to before, my results shed light on the impact of the change brought by Bayh-Dole.

A natural additional test stems from the fact that, because Bayh-Dole affected federally funded inventions, universities receiving more federal funding before the law was passed were effectively more "treated" by the change: they had a larger affected research base from which the local economy could now benefit. I thus test whether areas grew differentially depending on the amount of federal research funding their local universities attracted in the several years before the Act was passed. I use detailed information on federal funding by agency and university to measure whether this effect holds for technological areas that might be especially closely tied to industry, such as those funded by the Department of Defense and the National Institutes of Health.

The Census Bureau's Longitudinal Business Database (LBD) enables me to measure outcomes for detailed industries at a high level of geographic specificity from 1977 to 1997. The detail of the data permits tight connections between university strengths and related industry employment, payroll, and establishment dynamics. Spanning twenty years, the data cover the passing of the Bayh-Dole Act and facilitate measurement of its long-run effects. This long horizon is crucial for understanding the substantial shift, marked by this policy change, in the relationship between research universities and their industrial neighbors.

I find that employment, payroll, and average wages grow differentially faster after the Bayh-Dole Act in industries more closely related to the technological strengths of nearby universities. The magnitudes – 31 employees and \$1.5 million in payroll per county-industry (100 worker base) after Bayh-Dole for a standard deviation increase in industry intensity, or 13 employees and \$665 thousand in payroll per effective patent – are considerable and grow with geographical proximity to the university, supporting the importance of spatial relationships in the spread of knowledge. Areas surrounding universities that received more federal research

funding before the law was passed grow faster after the law than do others; the effect is particularly large for DOD and NIH funding.

Isolating the effects of university innovation on local industry is a fundamentally difficult task because universities have developed together with their local economies over time, influencing each other and each being influenced by similar area fundamentals. University and industrial activity are thus naturally correlated: communications technologies, for example, developed in firms like Federal Telegraph in the nascent Silicon Valley just as they were developing in Stanford University laboratories. Use of a national external shock to this system, like Bayh-Dole, and in particular one whose theoretical impact differs across geographical and technological areas, brings new identification to this long-standing measurement problem.<sup>4</sup> This strategy assumes that the industries and places I measure to grow the fastest after Bayh-Dole weren't already on faster growth trajectories, for reasons other than university research, before the law. It is possible, however, that there were pre-existing industry growth trends that continued after Bayh-Dole and explain my effect; these cannot be measured in the LBD data because they do not stretch back far enough.<sup>5</sup>

Having shown significant employment and payroll growth effects, I investigate which types of establishments – entrant versus incumbent, large versus small – are most complementary with university innovation in producing employment growth. Do universities generate growth via entrepreneurs or existing powerhouses?

---

<sup>4</sup> Previous research has investigated the effects of university spending, as opposed to innovation, on growth. For example, Aghion et. al. (2009) use political instruments to identify the effects of state university spending on growth. Kantor and Whalley (2010) use changes in endowment spending to measure effects of university activity on local labor income.

<sup>5</sup> Data on subsets of my sample do exist far enough back to be able to address this concern to some extent. While industry-level County Business Patterns data are suppressed too severely to be of much use, confidential data from the Census of Manufactures could be used to estimate pre-trends for industries in the manufacturing sector.

While entrants may be highly innovative and have been shown to drive year-to-year employment growth economy-wide (Haltiwanger et. al. 2010), incumbents may have established relationships with universities, distribution channels, manufacturing expertise, and brand name that help to preserve their power. I find that entering establishments, rather than incumbents, drive the twenty year growth effect from universities. Multi-unit expansion rather than single-unit entry accounts for 70% of employment growth from entrants, indicating the smaller role of entrepreneurs.

Some theories of innovation and growth further predict that innovating entrants may spur technologically advanced incumbents to compete while forcing laggards into decline, generating a bifurcated effect (Aghion et. al. (2005); Aghion et. al. (2009)). I find evidence consistent with this pattern and suggestive of creative destruction: incumbents experience more turnover in industries more closely tied to university innovation. Meanwhile, select incumbents in those industries appear to remain competitive and grow.

Like entrants, small firms – and small university spinoffs in particular – are often considered to be drivers of growth (Shane (2004)).<sup>6</sup> But there may be economies of scope (Arrow (1962)) or scale (Chandler (1990)) in innovation – or in translating innovation into growth via mass production and marketing of new products – that might lead large firms to be complementary to university research. I thus decompose the university growth effect into entry size categories and find that the largest establishments contribute considerably more absolute employment than do small ones. Small establishments, however, enter in larger numbers in university-tied industries and contribute substantial employment growth in proportion to initial size, lending some support to the Shane view of the importance of spinoffs.

---

<sup>6</sup> We know much, descriptively, about differences in spinoff activity across institutions and industries but less about spinoffs' systematic contribution to growth.



University effects are relevant for policy on multiple levels. National intellectual property policy, such as that contained in the Bayh-Dole Act, and federal research subsidies to universities both aim to enhance economic benefits derived from universities.<sup>7</sup> Local policy makers in particular care about the benefits that accrue to university areas. Though many public attempts to generate clusters of innovation fail at huge cost (Lerner (2009)), other evidence suggests that skilled populations play an important role in subsequent area income (Glaeser et. al. (1995); Moretti (2004)). My results emphasize the importance of universities for local growth.<sup>8</sup>

### **Health Insurance Costs and Young Firms**

Medical costs in the United States have been rising, and with them health insurance premia. According to a 2010 Kaiser Family Foundation national survey of firms, average premia have increased as much as 114% since 2000, though plan generosity has generally declined or remained constant.<sup>9</sup> The response to such changes is often centered on concerns of affordability and the uninsured, but because of the tying of insurance to employers, there may also be effects of rising medical costs on firms. Indeed, "health care is the number one concern for small businesses today," says Karen G. Mills, spokeswoman for the Small Business Administration (SBA). "The current health system places a heavy burden on many small business owners

---

<sup>7</sup> Because ideas are a non-rival good that can generate increasing returns to scale, they may be privately under-produced and warrant public subsidies and/or intellectual property protection to increase the expected return to investing in discovery (Romer (1986); Jones and Romer (2009)). The U.S. government engages in both of these means of encouraging research.

<sup>8</sup> Creating a strong university out of nothing is obviously a difficult task. Using universities as a policy tool seems more reasonable in the context of marginal research subsidies or intellectual property law. Furthermore, whether a policy to stimulate a local economy via its university would be welfare-enhancing nationally is a separate question (Glaeser and Gottlieb (2008)). Place-making policies require more justification than just evidence of stimulative local effects. In particular, they require non-linearities in agglomeration economies: benefits of one place growing outweighing losses from another shrinking. In the case of universities and local industry growth, it may well be that firms shift operations away from other geographical areas to be close to and benefit from universities. I do not address net effects of universities on industry growth across geographic areas.

<sup>9</sup> Kaiser Family Foundation and Health Research and Educational Trust, "Employer Health Benefits: 2010 Summary of Findings" (2010).

because it does not provide access to affordable coverage for them and their employees."<sup>10</sup>

Policy-makers have acknowledged the difficulties small firms may face and have exempted firms with fewer than 50 employees from the employer mandate specified in the recent health care reform.<sup>11</sup> Meanwhile, economic theory suggests it's not obvious – though certainly possible – that small firms will be negatively affected by the rising cost of a benefit. Substantial work has investigated effects of benefit costs at the individual level, but less is known about how rising health costs affect the viability of small firms. This relationship is difficult to measure because of unobserved local factors that may drive both premium growth and small business success. This paper uses state- and state-year-level legal changes that affect subsequent premium growth to identify the effects of rising health costs on young, small firms' survival, employment growth, and propensity to offer health insurance to employees.

States have enacted a number of legal measures in an attempt to either increase coverage or keep medical cost growth down. One class of regulation to increase coverage targets insurers, placing restrictions on their ability to reject applicants and on the range of prices they are allowed to charge. In particular, community rating laws limit the extent to which insurers can adjust premia for the health status and/or risk factors of an individual or small group. Though the goal was to increase coverage, in practice the price compression across types led healthy types to drop out of the market and premia to spiral upward in states that enacted such laws (Clemens (2011); Herring and Pauly (2006)).

Another class of state legal changes reformed the tort system in multiple ways, from altering liability rules to capping damages to regulating contingent payment contracts of lawyers.

---

<sup>10</sup> "Health Care Reform is Crucial for Small Businesses," Op-ed, Office of the Administrator, Small Business Administration, 7/30/2009.

<sup>11</sup> Employer requirements are set by the Patient Protection and Affordable Care Act of 2010 to go into effect in 2014.

Although the effect on medical costs of several of the laws is theoretically ambiguous, a number of tort reform laws have been shown in practice to reduce health insurance premia in the states in which they were enacted (Avraham et. al. (2009)). The staggered state enactments and strikedowns of caps on non-economic damages and punitive evidence reform will help to identify the effects of premium growth on small firms.

My strategy employs community rating laws and each of these two tort reforms separately to instrument for premium growth, generating three independent estimates of the effects of rising health costs on small businesses. I employ data from the restricted-access Kauffman Firm Survey (KFS), which enables me to follow a set of new, independent firms - a category we know from Haltiwanger et. al. (2010) to be a particularly important one for job creation. Because the KFS is the largest existing longitudinal survey of new firms and contains information on owners and health insurance decisions, it is quite appropriate for studying the interaction of two focuses of public policy: addressing rising health costs and encouraging entrepreneurship. Using state-employer size-year level data on insurance premia from the Medical Expenditure Panel Survey (MEPS) avoids the dual issues of (1) endogeneity of firm-level premia to employees' health status and risk factors, and (2) the need to impute premia for firms that don't offer insurance.

I find that the survival and employment growth of new firms faced with faster premium growth are not significantly affected, and I can rule out large negative effects. Perhaps more surprisingly, new firms faced with faster premium growth are significantly more likely to offer health insurance, contrary to what one might expect given previous evidence indicating a negative elasticity of health insurance demand of the self-employed with respect to its after-tax price (Gruber and Poterba (1994); Heim and Lurie (2009)). However, if the variation I use in premium growth is actually picking up a difference in underlying medical costs – rather than a

pure price difference as in many previous studies – the result may reflect employees' increased value of insurance as the cost of the bad state of the world has gone up, combined with the tax advantage of purchasing insurance through a firm.

That firms should be affected by benefit provision is not to be taken for granted. Summers (1989) shows that, in a simple model of benefit provision, wages should decline to offset exactly the cost of the benefit if the employee values the benefit at its cost to the firm. But wages may not be able to adjust perfectly – for instance because of downward nominal wage stickiness – leaving some of the incidence on the employer. Small firms also may be unable to adjust wages fully without losing high quality labor to larger firms, which in many cases are able to provide health benefits at lower administrative loads and lower cost (Helms et. al. (1992); CBO (2000); RWJF (2002)).

While economists have paid much attention to the effects of health insurance prices on the transition to self-employment (Holtz-Eakin et. al. (1996)),<sup>12</sup> take-up by the self-employed (Gruber and Poterba (1994); Heim and Lurie (2009)), take-up by the employed conditional on availability through the firm (Gruber and Washington (2005)), take-up by individuals (Cutler et. al. 2005), and individual-level wages, hours worked, and part time/ full time employment (Cutler and Madrian (1996); Baicker and Chandra (2006)), less is known directly about the effects on small (employer) firms.<sup>13</sup> Nevertheless, this is a group that represents a large share of employment and whose advocates claim it is affected substantially by rising health costs.

---

<sup>12</sup> See also Fairlie et. al. (2011) for a study of the effects of a potential disruption in health insurance, coupled with health status, on propensity to become self-employed.

<sup>13</sup> Some small-sample panel studies have been done looking at small firms' decision to offer insurance as a function of its price, but these tend either to impute premia for firms that don't offer insurance (Feldman et. al. (1997)) using biased models, or to rely on demonstration projects in which only a limited set of insurance packages were offered to firms who then took them up at low rates (Helms et. al. (1992)).

Moreover, the prevalence of small businesses in an area strongly predicts subsequent growth (Glaeser et. al. (2010)).

The result that young firms are not significantly negatively affected by health costs on the margins of survival, employment, and insurance availability does not rule out the possibility that these firms are adjusting on other margins, such as the quality of labor. But these results provide no compelling evidence that rising health costs are substantially detrimental to young firms.

### **Capital Gains Taxes and Household Portfolios (with Stephen Miran)**

Capital gains taxes, paid by households when they sell an asset that has appreciated in value since purchase, are a subject of significant policy interest. Tax rates have been adjusted several times in the past few decades and further changes are often under consideration. Now, as current tax rates are scheduled to “sunset” in 2011 to the higher pre-2003 rates, the subject is as relevant as ever; President Obama campaigned in part on an increase in the capital gains tax rate faced by high income earners.

Aside from partisan views as to how much the owners of capital should be taxed, the magnitude of the tax is important because of the behavioral distortions and resulting deadweight loss it may inflict on the economy. Because the tax may induce asset holders to sell assets less frequently than they optimally would – an effect known as “lock-in” – individuals may hold the wrong set of assets and leave capital in the wrong firms. There may be further distortionary effects as people accommodate increased friction in asset sales with substitution between asset classes, between labor and capital income, or with changes in the timing of their consumption streams.

We use a detailed data set on households' investment accounts at a large discount brokerage firm and employ several strategies to investigate the role of capital gains taxes in households' investment behavior. First, we estimate a reduced form model of within-account asset rebalancing to adjust portfolio weights. An investor who does no trading from one month to the next will experience some “passive change” in the composition of his portfolio, as the market causes some assets to appreciate and some to depreciate. If the investor is rational and there are no frictions, he should actively rebalance his portfolio to undo some of the passive changes. The extent to which he counters passive changes with active changes – transactions that bring him back towards his optimal share in risky assets – has been shown by Calvet, Campbell, and Sodini (2009) to be both incomplete and related to a number of demographic factors likely correlated with financial sophistication. Individuals, they find, tend to offset on average about half of the idiosyncratic passive variations they experience in their risky share. Because capital gains taxes act as a proportional transactions cost on appreciated asset sales, they theoretically increase the cost of portfolio rebalancing and could potentially induce people to rebalance less. We thus extend CCS's estimation of rebalancing by incorporating taxes: how much less do households rebalance as a result of capital gains taxes?

Because we observe households' positions and transactions separately in each of their taxable and tax-deferred (IRA and Keogh accounts; henceforth “TDA”) investment accounts, we can compare behavior across account type within households to measure the effect of the tax. Our strategy amounts to using the tax-deferred account as a counterfactual where households reveal their preferences in the absence of taxes.<sup>14</sup> We find that there is approximately one-fourth less rebalancing in taxable accounts.

---

<sup>14</sup> Since our comparison is across accounts, we do not attempt to explain why investors make suboptimal portfolio decisions to begin with.

Since gains realizations can be made tax-free in retirement accounts, rational and attentive investors should, to some extent, trade in their retirement accounts to offset passive changes in their taxable investment accounts. We thus measure the extent to which households may strategically use tax-free accounts in rebalancing. This test is interesting not only for its implications for the capital gains tax but also as a test of mental accounting, a much discussed phenomenon in the behavioral economics literature that has nevertheless rarely been given a rigorous test. The extensive literature on mental accounting (see the review in Thaler, 1999) tends either to study general equilibrium asset pricing implications of mental accounting or to look for indirect evidence of the phenomenon in contexts like work hours or calorie consumption.<sup>15</sup> In contrast, we attempt to provide a more direct test of mental accounting by measuring whether agents offset passive changes in one account by trading in another account, a necessary condition for optimization of the unified portfolio in the presence of taxes. Our results indicate a lack of cross-account rebalancing and a failure for households to behave in a tax-minimizing manner, which we interpret as indicative of mental accounting.

Having estimated reduced-form differences in rebalancing across account types, we seek to understand this behavior in more detail. Because theoretical literature on optimal control in an environment with transaction costs tends to find the optimality of  $(S,s)$  adjustment processes, we attempt to fit a model of these lumpy adjustment rules. We measure the triggers and targets of an  $(S,s)$  band around households' optimal risky asset share under the hypothesis that the taxes may widen the band. When we force all households into an  $(S,s)$  model, the results do not conform to the economic logic of lumpy adjustment in this setting: two of the band's estimated

---

<sup>15</sup> Indirect evidence of mental accounting in work hours might be something like seeing workers stop at a target number of hours per day, even if the wage may be higher that day and thus warrant substitution from another day or week. A recent exception which better tests mental accounting is Choi, Laibson and Madrian (2008). Our study differs in examining mental accounting in portfolio rebalancing.

parameters are out of order. We therefore turn to a mixture model, where some agents follow this “rational model,” and others follow a “behavioral” model of churning. The mixture model is estimated via expectation maximization (EM).<sup>16</sup> When we allow for a portion of the households to be behavioral, the estimated (S,s) band has appropriately ordered parameters and is indeed wider in the taxable account, as predicted by a model of lumpy adjustment with transactions costs. However, EM assigns less than 10% of the households to the rational model. We find substantial heterogeneity among our households, mostly consistent with previous literature: households more likely to be behavioral are those which are more poorly diversified, trade more, and have lower income.

Having established the existence and magnitude of the distortion in trading behavior induced by the tax, we produce estimates of the resulting welfare loss to individual investors. Instead of relying on a particular theory of behavior to set a benchmark of optimality, our approach uses the observed behavior in the TDA and interprets it as the revealed preferred behavior in an untaxed setting. The strategy consists of estimating, for each household, the empirical distribution of returns it experiences. With this distribution, and allowing only returns to differ across accounts, we can simulate the expected utility from end-of-horizon consumption of the assets in each account. The utility cost of having the “wrong” asset mix consists of the utility difference between a household's accounts, when the only thing allowed to differ is the asset mix and the tax. Our analysis suggests that, for plausible degrees of risk aversion, the median household bears a deadweight burden of 10-25% of the revenue collected. We find substantial heterogeneity of welfare costs. Households that are better diversified, have more wealth in these accounts, sell less, and are less likely to churn suffer lower welfare costs; we interpret these households to be more sophisticated and better able to avoid large tax burdens.

---

<sup>16</sup> EM is discussed in Section 3.5. It is a statistical method for estimating mixture models.



Our study of the dynamic portfolio problem focuses on how investors maintain their risk-return profiles. In contrast, much previous empirical research on distortion from capital gains taxes has focused on the elasticity of gains realizations with respect to the marginal tax rate, or on asset location, choice, and sales. While some of these parameters are useful for tax revenue predictions, their implications for economic well-being are unclear; the focus on realizations and asset location ignores the overall asset mix, and thereby fails to separate the consumption and portfolio problems. Our goal in this paper is to shift the focus on capital gains taxes from effects on sales/outflows to effects on risk maintenance. In doing so we isolate the portfolio problem from the consumption problem and attempt to link the literature on capital gains taxes with the growing literature on household finance, especially with work by Calvet, Campbell and Sodini (2007, 2009).<sup>17</sup>

Since the seminal work of Feldstein, Slemrod and Yitzhaki (1979), estimates of the elasticity of gains realizations have undergone many iterations; for instance, Burman and Randolph (1994) attempt to purge the estimation of intertemporal substitution, while Auerbach and Siegel (2000) allow for investor heterogeneity. More recent papers have edged the literature away from the elasticity parameter by asking how other financial behaviors interact with taxes. Ivkovich, Poterba, and Weisbenner (2005) estimate how the probability of selling an asset is affected by capital gains taxation. They compare hazard rates in people's taxable accounts to those in their tax-deferred (retirement) accounts and attribute the difference to the tax. Bergstresser and Poterba (2004) and Poterba and Samwick (2003) study the relationships between marginal tax rates, asset location, asset choice and tax efficiency.

Our approach is similar to that of Poterba, Ivkovich and Weisbenner, but we study the portfolio as a whole, whereas they study sales of individual assets. As Auerbach and Siegel

---

<sup>17</sup> Campbell (2006) provides an excellent introduction to the literature on household finance.

emphasize, there is a wide gap between the existing empirical literature and the theoretical focus on tax avoidance in Balcer and Judd (1987), Constantinides (1984) and Stiglitz (1983). By focusing on state variables, we attempt to bring the empirics marginally closer to the theory of portfolio choice and consumption behavior, and closer to the theory of behavioral finance. Doing so allows us to attempt an estimate of investors' utility losses that result from lock-in.