

Taxation and the Allocation of Talent*

Benjamin B. Lockwood[†] Charles G. Nathanson[‡] E. Glen Weyl[§]

December 2012

Abstract

Taxation affects the allocation of talented individuals across industries by blunting material incentives and thus relatively magnifying the non-pecuniary benefits of pursuing a “calling”. If higher-paying industries (e.g. finance and management) generate less positive net externalities than lower-paying professions (e.g. public service and education) this may enhance efficiency. We develop a theory of income taxation as implicit Pigouvian taxation of these externalities and calibrate it using data on the distribution of income and talent across industries. Even without any redistributive motive, tax rates are highly sensitive to the externalities assumed within a spectrum many would consider reasonable: they range from extremely regressive to highly progressive at high incomes. Our theory thus offers an alternative, pure efficiency rationale for non-linear income taxation, challenging the connection between high long-run labor supply elasticities and low optimal tax rates and motivating further study of the externalities generated by professions.

*This paper builds on and replaces a 2007 working paper solely authored by Weyl, “Psychic Income, Taxes and the Allocation of Talent”. We gratefully acknowledge the financial support of the Alfred P. Sloan and Marion Ewing Kauffman Foundations. We owe a debt to Claudia Goldin for providing access to data. We also appreciate the helpful comments of Tony Atkinson, Raj Chetty, Andrei Shleifer and Joel Slemrod. We are especially grateful to Joshua Bosshardt for writing the applet and to Daichi Ueda for more general research assistance. All errors are our own.

[†]Department of Economics, Harvard University: Littauer Center, 1805 Cambridge Street; Cambridge, MA 02138: lockwood@fas.harvard.edu.

[‡]Department of Economics, Harvard University: Littauer Center, 1805 Cambridge Street; Cambridge, MA 02138: cnathans@fas.harvard.edu.

[§]Department of Economics, University of Chicago: 1126 E. 59th Street, Chicago, IL 60637: weyl@uchicago.edu, <http://www.glenweyl.com>.

If we don't have an economy built on bubbles and financial speculation, our best and brightest won't all gravitate towards careers in banking and finance. Because if we want an economy that's built to last, we need more of those young people in science and engineering. This country should not be known for bad debt and phony profits. We should be known for creating and selling products all around the world...

- President Barack Obama, Speech at Osawatomie High School, December 6, 2011

1 Introduction

The allocation of talented individuals across professions varies tremendously across time and space. For example, according to data collected by Goldin et al. (2013), two to three times as many male Harvard alumni from the 1969-1972 cohorts pursued careers in each of academia and non-financial management than pursued careers in finance. Twenty years later, careers in finance were fifty percent more common than in academia and comparable with those in non-financial management. If private product is anywhere near social product, these talented individuals constitute a large fraction of many societies' human capital: in the United States, for example, roughly half of all income is generated by the top 10% of income earners and nearly a quarter is generated by the top 1% (Atkinson et al., 2011). If, as Baumol (1990) and Murphy et al. (1991) argue, different professions have different ratios of social to private product (viz. some have negative and others positive externalities) then these differences in talent allocation across societies may have important implications for aggregate production.¹ In this paper we argue that progressive income taxation is a powerful tool affecting the allocation of talent and therefore that, if the above logic is correct, the effect of such taxes on the allocation of talent is central to their optimal design.

Our argument is that in selecting an industry, talented individuals face a trade-off between pursuing a “calling” that offers them high non-pecuniary benefits and choosing a career that offers better remuneration. Higher marginal tax rates between the income earned in the lower-paying and higher-paying career make the latter relatively less attractive by narrowing the material sacrifice associated with following a passion. To the extent that better paying professions are also more likely to generate negative (less likely to generate positive) externalities, raising marginal tax rates has a pure efficiency benefit. Conversely, if one believes that better paying professions actually generate more positive externalities, efficiency might require negative marginal tax rates for the well-off.

Income taxation is admittedly a blunt tool in addressing differences in externalities across professions. We nonetheless believe that analyzing the effects of non-discriminatory income taxation is

¹This argument has drawn significant public attention in recent years. For example, it sparked the foundation of the “Stop the (Wall Street) Brain Drain” movement.

useful for policy for several reasons.² First, occupations are sufficiently difficult to define objectively that any career-specific tax scheme would be subject to manipulation that would undermine its efficacy. Second, the political economy consequences of allowing profession-specific taxation could be dangerous, unleashing a range of special interest lobbying and propaganda that is unlikely to lead to an efficiency-improving equilibrium. Third, considerations of horizontal equity may make differential taxation of different occupations ethically or politically unpalatable. Finally, profession-specific taxation is simply not on the public agenda, while income tax reform is; thus, given the current second best situation, we believe economists' views of optimal taxation should be influenced by their goals in allocating talent, not just their views about redistribution. Many of these concerns are analogous to those that led Mirrlees (1971) to focus attention on non-linear income taxation and assume that wages as such were non-contractible even though they seem to be at least partially observable in practice. In Subsection 5.5 we will, in a future draft, discuss how much better our work suggests profession-specific taxation could do compared to optimal non-linear, horizontally equitable taxation. Rothschild and Scheuer (2012), and to a lesser extent Philippon (2010), take a similar approach to ours, though our interest is in quantitative calibration in contrast to their focus on directional theoretical results; we discuss the relationship between these papers, which were developed independently of ours several years after the first 2007 draft of this paper was circulated, more extensively in Subsection 4.3.

It may be useful to briefly contrast our theory of optimal income taxation to the classical approach based on the work of Mirrlees and survey recently by Mankiw et al. (2009) and Diamond and Saez (2011). Our theory requires no distributive motive to generate high marginal tax rates; even when the social planner seeks only to maximize total social wealth, very high and sometimes highly progressive tax rates may be justified. For example, we find that for assumptions about externalities matching the views of two of the authors (Lockwood and Weyl), marginal tax rates over some ranges are significant (25-50%) over a wide range; for externality assumptions that we see matching the views of the "Occupy Wall Street" movement, pure efficiency-maximizing marginal rates are often near 50%. Elasticities, intensively of labor supply or extensively of career switching, play at most a secondary role in our theory, just as they play no role in Pigou (1920)'s theory of taxation. The primary aim of policy is to implicitly target taxation at professions generating negative externalities and subsidies at those generating positive externalities, thereby to shifting individuals to the most productive professions rather than to extract maximal revenue while causing minimum shifts in economic activity. In that sense it more closely resembles the "Just Desserts Theory", advanced by Mankiw (2010), that taxation should be based on ensuring individuals receive their social contribution than it does the veil of ignorance logic of Harsanyi (1953) and Rawls (1971)

²In complementary work, Posner and Weyl (2013a) and Posner and Weyl (Forthcoming) discuss more detailed policies for correcting the externalities in two of the professions that many believe negative externalities are strong (finance and law, respectively). We believe that both targeted, but potentially bureaucratic and manipulable, and blunter but more robust and simpler policy interventions have a valuable role to play.

on which the Mirrlees theory is founded.

Finally, and more stylistically, we do not seek, as many recent papers on optimal taxation have, to precisely estimate, or even calibrate, optimal tax rates. The central parameters in our theory, the degree to which different occupations generate externalities, are simply too uncertain and controversial for the results to be broadly plausible. Rather, we seek to persuade the reader that optimal tax policy towards the top part of the income spectrum may be extremely sensitive to estimates of these parameters over currently reasonable ranges of disagreement and uncertainty, more sensitive than they are to the traditional elasticity parameters on which the literature has heretofore focused. For example, while we find high marginal rates are optimal for “left-wing” assumptions about externalities, when we choose values to mirror our guess of the views of the “Tea Party” movement, marginal rates are always significantly negative, particularly on the wealthy, sometimes reaching a subsidy of 50% of income. But we do not simply consider several fairly arbitrary cases; instead we include an applet, written jointly with Joshua Bosshardt, at <http://home.uchicago.edu/weyl/taxapplet.html> that allows the reader to input her desired parameter values and reads off optimal tax rates for any given views she has about externalities. We hope that this observation will stimulate research on estimating these parameters more precisely.

Outline

We begin in Section 2 by briefly summarizing the data we use to calibrate the model in the United States. The central information our theory requires is the distribution of earnings and talented individuals across industries. We draw the first from data on earnings distributions available from the Bureau of Labor Statistics, supplemented by data on the top end of the income distribution drawn from Bakija et al. (2012).³ We draw the second from data collected by Goldin et al. and referenced above on the career choices of alumni from Harvard University, comparing this with the distribution from the 1950’s to gauge labor supply responses.

In Section 3, we consider the simplest setting for our analysis: one in which there are no income effects or redistributive motives and the planner simply seeks to maximize total social income (Kaldor-Hicks efficiency). After setting up a model and providing a general characterization of optimal taxes in this setting in 3.1, we further simplify the analysis by assuming that, conditional on income levels, there is no correlation between externalities and either intensive elasticities or career switching into or out of that income level. Then we show, in Subsection 3.2 that optimal taxation simply requires equating tax at a given income level to the average externality created by individuals who earn that income level. If the only relevant elasticity is career switching then average tax rates are matched to (average-given-income) externalities, a policy we call “Average Tax Externality Matching” (ATEM). If the only relevant elasticity is intensive, on the other hand, marginal tax

³In a future draft we hope to further supplement this, at the very highest end, with the data of Kaplan and Rauh (2010). However for the current draft we ran into technical difficulties with interpreting these data.

rates are matched to externalities, a policy we call “Marginal Tax Externality Matching” (MTEM). Note that this optimal tax scheme thus depends mostly on the externalities created in each profession and only on elasticities (intensive v. extensive) to the extent that they modulate whether MTEM or ATEM is optimal.

This makes optimal taxes particularly easy to calibrate using the data described in Section 2; each of MTEM and ATEM can be calculate easily and compared. In Subsection 3.3 we show the results of these calibrations for five assumed assignments of externalities to professions that correspond to our subjective guess about the views of various parts of the American political spectrum, as well as to views to which different co-authors of this paper are sympathetic. The results are fairly robust to the comparison of ATEM and MTEM, though a bit more extreme under ATEM. In both cases, marginal taxes are modest (below 20%) at most incomes for the views favored by the authors, either because (Nathanson) they are skeptical of many externalities existing or because (Lockwood and Weyl) the professions we believe have large externalities (finance and law negative, teaching, academia and entrepreneurship positive) are either a small fraction of individuals at every income or balance one another conditional on income.⁴ On the other hand, for views we impute to correspond to the Tea Party (government and media have negative externalities, business and finance have positive externalities) and Occupy Wall Street (essentially the reverse of the Tea Party) movements, marginal tax rates are as discussed above. We plan to use our applet described above to collect data on clusters of views about externalities that we will use in future drafts of the paper.

In Section 4 we discuss the relationship of our theory to several literatures, which we largely neglect until that section. First, in Subsection 4.1 we challenge the standard argument (Keane, 2011) that high long-term labor supply elasticities imply that low marginal tax rates on high earners are optimal; in fact, the reverse may often be the case. Regardless, in our theory the externalities created by different occupations, which have received limited attention, rather than labor supply elasticities, which has been the subject of extensive study, are the key parameters of empirical interest for tax policy. Second, in Subsection 4.2 we argue that our framework offers a natural context for empirically calibrating Holmström and Milgrom (1991)’s theory of multi-tasking in agency relationships, for the first time that we are aware of. Third, income taxation is a straightforward policy lever for shifting the allocation of talent, the literature which motivates our work. Under Lockwood and Weyl’s externality shares, the shift in the allocation of talent implied by the Goldin et al. data and largely reversible by changing tax rates could account for as much as half of the increase in top income shares over the last half century or around one sixth of the decline in GDP growth from the 1945-70 period compared to the 1980-2005 period. Finally, we use anecdotal and public opinion evidence to argue that many of the debates, both in the public

⁴However, Lockwood and Weyl’s views call for large subsidies for those earning middle incomes (\approx \$70-130k) that are fully phased out before they reach the rich (\approx \$250k).

political sphere and in non-Neo-Classical economics, over optimal taxation are driven by the factors we emphasize (viz. the externalities of different occupations) rather than the factors emphasized by the Mirrlees model (viz. distaste for inequality and the responsiveness of work to incentives).

In Section 5 we consider several extensions of our basic model to make it richer and more realistic and to draw out its implications more fully. By assuming random substitution of individuals into and out of income levels our basic analysis may bias marginal rates towards zero. Thus in Subsection 5.1 we consider a model where individuals have a general ability that determines their position in the profession-specific income distribution. We show analytically that this makes even more extreme the effects of externalities on optimal marginal tax rates when there are two professions, one of whose income distribution stochastically dominates the others. In Subsection 5.2 we calibrate a standard structural choice model under the general ability assumption and show that, as suggested by our theoretical results, marginal rates are much more extreme than under either ATEM or MTEM: for example, under Lockwood and Weyl’s profiles they are similar to the Occupy results under ATEM.

To contrast our framework with the standard Mirrlees approach, in Subsections 5.3 and 5.4 we plan, in a future draft, to add a redistributive motive to the model and compare the relative importance of externalities compared to other standard parameters to determining optimal tax rates. Subsection 5.5 will, in a future draft, quantify various welfare statistics: the gains from profession-specific v. horizontally-equitable taxation, the gains from moving to optimal taxes from the current levels, etc. Because most calibrations that call for high taxes require low elasticities, welfare gains are typically modest and income gains are negative. It seems plausible, by contrast, that given the relatively large elasticities we find, some of our assumptions about externalities would generate large gains to both welfare and income from higher marginal tax rates.

We conclude in Section 6 by discussing how future research might clarify some of the crucial parameters on which we argue optimal income taxation depends.

2 Data

The three central inputs to our analysis are 1) the distributions of income within different professions, 2) the distribution of individuals across professions (and how this responds to changes in wages) and 3) the externalities of different professions. In this section we discuss how we obtain 1) and 2). Because of the absence of data on 3) we treat these parameters as basically subjective for the moment, as we discuss in Subsection 3.3 below.

To calibrate 1), profession-specific income distributions, we combined data from the Bureau of Labor Statistics (BLS) with data published by Bakija et al. (2012). Bakija et al. group the industries of earners in the top 1% into 13 categories plus no occupation and other. These are, with their accompanying abbreviations, academia/research (ProfSci), advertising/market/sales (sales), agriculture (farmer), arts/entertainment (artsEntSport), business operations (busOp, primarily man-

agement consultants and accountants), engineering/technical (compMathEng), entrepreneurship (classified for now in “other” for technical reasons; to be corrected on a future draft, but small), financial services (finance), law (law), management outside of finance (management), medicine (medicine), public service/military (government) and real estate (real estate). Bakija et al. obtain these, in a manner described in greater detail in their paper, by merging together various North American Industry Classification System (NAICS) and Standard Occupational Classification (SOC) code pairs. In a few cases we were not able to follow their classification exactly and we also add a category for primary and secondary education (teacher). In Appendix A we describe our procedures for constructing these classifications where they do not match Bakija et al. exactly.

The BLS publishes the 10th, 25th, 50th, 75th and 90th percentiles of NAICS-SOC code pair-conditional income. We merge together NAICS-SOC pairs as described above; because the BLS also publishes data on the number of individuals in each code, we can obtain the distribution of the merged occupation if we recover a income distribution conditional on NAICS-SOC pair. We do this by fitting a log-normal distribution, which is known to fit income distributions well below the 90th percentile, to the five BLS provided data points in each NAICS-SOC pair. The fit is very close in essentially all cases. We use this to construct the income at the 10th, 25th, 50th, 75th and 90th percentile of each profession as defined by Bakija et al.⁵ We then obtain two additional data points by using the Bakija et al. data, which reports the number of individuals in each profession earning above the cut-offs for the 99th and 99.9th percentiles of the income distribution (which they numerically report).

Thus we obtain a semi-parametric estimate of five points in the lower 90 percentiles of each profession-specific income distribution and a non-parametric observation of two points at the upper tail of the cumulative distribution function of each profession-specific income distribution. To these 7 points we fit a log-normal Pareto distribution (a distribution that is log-normal below a threshold and Pareto above with the parameters and the transition point estimated by maximum likelihood). The transition to Pareto occurs in the top 10 percentiles of each profession-specific distribution (except in the primary/secondary education profession where the Bakija et al. data is not reported and thus we just use a log-normal distribution). Further details of these estimations are reported in Appendix A.

Using these estimated distributions and the BLS’s data on the total number of individuals in each of these professions we obtain the income-conditional distribution of individuals across professions, the first part of 2) and the primary input to our calibration in Subsection 3.3. To illustrate this key input, Figure 1 shows the fraction of individuals in various professions conditional on income levels ranging from \$10k to \$1 million. The x-axis is spaced logarithmically, with \$100k being the midpoint, but the labels on the x-axis are misleading in the current draft. The salient feature, which drives our results, is that nearly every profession outside of management, medicine and

⁵We plan to use these data in a more comprehensive way in a future draft.




Figure 1: Shares of individuals in various professions, conditional on annual income levels ranging from \$10k to \$1 million. X-axis is logarithmic, so the mid-point is \$100k, but the labeling of the x-axis is incorrect. Professions, other than “teacher” correspond to those in Bakija et al. (2012).

finance is represented significantly only at middle incomes (roughly \$70k-\$130k). Medicine has significant representation among the modestly wealthy (\$150k-\$500k) and finance has a substantial representation only among the fairly wealthy. But the overwhelming majority of the wealthy are non-financial managers and this pattern grows stronger the higher income level one considers.

These results were surprising to us, given that Kaplan and Rauh (2010)’s data, and our anecdotal experience from our own lives, are suggestive of finance and law accounting for many more of the wealthy than does non-financial management. The data of Goldin et al. (2013) are also suggestive of a much larger role for finance and law in high income groups. We plan to investigate this issue further in a future draft, but in the mean time we report normative results based on these positive findings; a stronger representation of finance and law at top incomes would significantly strengthen our results as disagreements about externalities of these professions are the strongest.

Finally, because we only have access to a single cross-section of top incomes by profession, we used the data of Goldin et al. to track the movement across professions over time. Goldin et al. surveyed, in 2005, members of three cohorts of Harvard alumni, each cohort consisting of four classes: 1969-72, 1979-82 and 1989-92. Each class was asked to identify, among other things, the career they followed (grouped in ways roughly corresponding to the professional classifications we use above) and their annual income. This allows us to observe shifts in occupational patterns over

time. We also use data from Philippon and Reshef (2012) on the evolution of wages in the financial sector over time in ways explained below.

3 Income Taxation as Implicit Pigouvian Taxation

In this section we construct the simplest and most intuitive version of our theory by ruling out any redistributive motive and considering the optimization of a non-linear, horizontally-equitable income tax to sort talent across professions to maximize total social income (Kaldor-Hicks efficiency).

3.1 General theory

There are N professions, $p = 1, \dots, N$ and a mass 1 of talented individuals. Each individual i is characterized by a $2N$ -dimensional vectors $\theta_i = (\mathbf{w}_i, \psi_i)$, where w_{ip} represents the money wage individual i would earn if she chose profession p and ψ_{ip} represents a parameter characterizing the non-pecuniary or *psychic* income she would receive from working in this profession. These characteristics of individuals are distributed in the population according to a non-atomic and differentiable distribution function f with full support on a convex and open $\Theta \subseteq \mathbb{R}_{++}^{2N}$. Each individual can choose how many hours to work, h , at a utility cost $\phi(h; \psi_{ip})$ where $\frac{\partial \phi}{\partial h} > 0$ (work is costly), $\frac{\partial^2 \phi}{\partial h^2} > 0$ (the costs of work are convex) and $\frac{\partial \phi}{\partial \psi} < 0$ (ψ shifts down the costs of working/up the non-pecuniary benefits of work). Wages are assumed linear, so the individual i 's income in profession p given that she works h hours is $w_{ip}h$. We do not constrain hours to lie in a finite range as we interpret hours more broadly as effort and we do not assume ϕ need be positive (work may be enjoyable, on net, but the marginal cost of hours is also positive).

The government must finance a net expenditure of I through the use of a non-linear income tax under which an individual earning total income $y = wh$, regardless of the profession in which she earns this wage, pays a total tax $T(y)$. Thus, just as in Mirrlees (1971), we effectively assume that neither wage nor profession can be verified or that, as suggested by Diamond and Saez (2011), some horizontal equity concern prevents greater discrimination.⁶ We assume this tax must be differentiable and that $T' \leq 1$ to avoid individuals disposing of income. Each individual's utility is just the sum of her money and psychic incomes.

Each individual is assumed to have quasi-linear utility in money income and thus to earn net utility

$$w_p h - T(w_p h) - \phi(h; \psi_{ip})$$

when she works h hours in profession p . Given the assumed convexity of ϕ , so long as T is also convex each individual has a unique optimal $h_p^*(w_p, \psi_p)$ to work conditional on being in profession p and will only move h_p^* locally in response to small marginal tax changes. We let $u_p^*(w_p, \psi_p)$ be

⁶See our discussion in the introduction for a more detailed justification.

the value of utility at this optimum. Sometimes we write $h_p^*(\theta)$, and similarly for u_p^* , which is interpreted as extracting the relevant components. Individual i chooses to work in the profession where her after-tax income plus her psychic income is highest, at her optimal profession-conditional hours level:

$$p^*(\theta) \equiv \operatorname{argmax}_{p \in 1, \dots, N} w_p h_p^*(w_p, \psi_p) - T(w_p h_p^*(w_p, \psi_p)) - \phi(h_p^*(w_p, \psi_p); \psi_p).^7 \quad (1)$$

Each profession has an externality share, e_p with the interpretation that an individual i working in profession p and earning wage w_{ip} generates a net externality on the rest of society, evenly distributed across individuals, of $e_p w_{ip}$ for each hour she works. The planner seeks to maximize total income (money net of externalities and psychic income) in her choice of the tax. In particular, she solves

$$\max_{T(\cdot)} \int_{\Theta} [(1 + e_{p^*(\theta)}) w_{p^*(\theta)} h_{p^*(\theta)}^*(\theta) - \phi(h_{p^*(\theta)}^*(\theta); \psi_{p^*(\theta)})] f(\theta) - I,$$

subject to the definition of h_p^* and p^* above. To derive optimal taxes, we follow the intuitive perturbation approach to the calculus of variations problem pioneered in economics by Wilson (1993) and in optimal income taxation by Saez (2001). Suppose the planner raises slightly the marginal tax rate at wage w , returning the raised revenue uniformly to the population so as to continue to satisfy her budget constraint and otherwise leaving fixed all other parts of the tax system. The redistribution thus induced has no net social value, as the planner seeks only to maximize total social wealth. Thus we can focus entirely on the behavioral responses to the tax rise.

One component of this is the local, *intensive* elasticity through the number of hours individuals choose to work. Given convexity of ϕ , so long as T is not too concave the optimal choice of h will always move locally in response to small changes in the optimal tax rate. In particular, if marginal taxes at income y , $T'(y)$, rises we can trace the impact on the optimal hours for an individual of type θ who is earning income y (because $w_{p^*(\theta)} h_{p^*(\theta)}^*(\theta) = y$), assuming she stays in the same profession, by the implicit function theorem. The first-order for $h_{p^*(\theta)}^*$ is

$$w_{p^*(\theta)} [1 - T'(w_{p^*(\theta)} h_{p^*(\theta)}^*(\theta))] = \phi'(h_{p^*(\theta)}^*(\theta); \psi_{p^*(\theta)}) . \quad (2)$$

We can now determine the effect of increasing $T'(y)$ by a small amount t using the implicit function theorem, letting $\epsilon_p^h(\theta)$ be the (negative) *intensive labor supply elasticity* of h with respect to the

⁷When the best profession for an individual is not unique, the tie may be broken in any manner as our assumptions on the distribution of wages and psychic incomes assure that the set of individuals facing such indifferences is of measure 0.

post-tax wage $w(1 - T')$:

$$-w_p \left(1 - \frac{\epsilon_p^h(\boldsymbol{\theta}) h_p^*(\boldsymbol{\theta}) T''(y)}{1 - T'(y)} \right) = -\frac{\epsilon_p^h(\boldsymbol{\theta}) h_p^*(\boldsymbol{\theta}) \phi''(h_p^*(\boldsymbol{\theta}); \psi_p)}{w_p [1 - T'(y)]} \implies$$

$$\epsilon_p^h(\boldsymbol{\theta}) = \frac{w_p^2 [1 - T'(y)]}{h_p^*(\boldsymbol{\theta}) [\phi''(h_p^*(\boldsymbol{\theta}); \psi_p) + w_p^2 T''(y)]}. \quad (3)$$

On the other hand, the marginal social value created by an individual working an additional hour is

$$(1 + e_{p^*}(\boldsymbol{\theta})) w_{p^*}(\boldsymbol{\theta}) - \phi'(h_{p^*}^*(\boldsymbol{\theta}); \psi_{p^*}(\boldsymbol{\theta})) = [e_{p^*}(\boldsymbol{\theta}) + T'(y)] w_{p^*}(\boldsymbol{\theta}),$$

where the equality follows by substituting in the first-order condition for hours, equation (2). Intuitively, by the envelope theorem, the net social value created by an additional hour of work is proportional to the private product (wage) multiplied by total externality associated with wages earned, both through the tax externality and the direct externality e .

If marginal tax rates rise at y , an individual of type $\boldsymbol{\theta}$ currently earning income y will change her hours by $\frac{\epsilon_{p^*}^h(\boldsymbol{\theta}) h_{p^*}^*(\boldsymbol{\theta})}{1 - T'(y)}$ and thus will change social welfare by

$$\frac{[e_{p^*}(\boldsymbol{\theta}) + T'(y)] \epsilon_{p^*}^h(\boldsymbol{\theta})}{1 - T'(y)}.$$

The set of all individuals earning income y is

$$\Theta(y; T) \equiv \{\boldsymbol{\theta} \in \Theta : w_{p^*}(\boldsymbol{\theta}) h_{p^*}^*(\boldsymbol{\theta}) = y\}.$$

Where it does not create ambiguity below, we drop the dependence on T . Let the density of such individuals be $f(y) \equiv \int_{\Theta(y)} f(\boldsymbol{\theta}) d\boldsymbol{\theta}$.⁸ The total impact on welfare through these local changes is then

$$\int_{\Theta(y)} \frac{[e_{p^*}(\boldsymbol{\theta}) + T'(y)] \epsilon_{p^*}^h(\boldsymbol{\theta})}{1 - T'(y)} f(\boldsymbol{\theta}) d\boldsymbol{\theta} =$$

$$\frac{f(y) \left([E[e_{p^*}(\boldsymbol{\theta}) | \Theta(y)] + T'(y)] E[\epsilon_{p^*}^h(\boldsymbol{\theta}) | \Theta(y)] + \text{Cov}(e_{p^*}(\boldsymbol{\theta}), \epsilon_{p^*}^h(\boldsymbol{\theta}) | \Theta(y)) \right)}{1 - T'(y)},$$

where the expectation and covariance operators are defined as usual, conditional on the relevant sets.

The second component of the behavioral response follows a similar normative logic, but is driven

⁸Note that this notation is a (valid) short-hand for this surface integral over a $2N - 1$ dimensional surface and should not be misinterpreted as a volume integral in the original space. For more technical discussion of the relevant mathematics see Veiga and Weyl (2012).

by changes in professions. In particular, let

$$\partial\Theta(y; T) \equiv \{\boldsymbol{\theta} \in \Theta : \exists p, q \in 1, \dots, N : (w_p h_p^*(\boldsymbol{\theta}) < Y < w_q h_q^*(\boldsymbol{\theta})) \wedge (u_p^*(\boldsymbol{\theta}) = u_q^*(\boldsymbol{\theta}))\}$$

be the set of *y-career switching* individuals who, under tax system T , are just indifferent between two professions, one of which has an optimal (for that individual) income level above and the other of which has an optimal income level below y . Let $f_S(y) \equiv \int_{\partial\Theta(y)} f(\boldsymbol{\theta}) d\boldsymbol{\theta}$ be the density of such individuals. Raising $T'(y)$ causes all of these individuals to switch from profession q to profession p .⁹ How does this change social wealth created?

In profession q , the individual i generates social wealth $(1 + e_{iq}) w_{iq} h_q^*(\boldsymbol{\theta}_i) - \phi(h_q^*(\boldsymbol{\theta}_i), \psi_{iq})$ while in profession p she generates social wealth $(1 + e_{ip}) w_{ip} h_p^*(\boldsymbol{\theta}_i) - \phi(h_p^*(\boldsymbol{\theta}_i), \psi_{ip})$. From the fact that she is indifferent between the two professions, we know that

$$w_{ip} h_p^*(\boldsymbol{\theta}_i) - T(w_{ip} h_p^*(\boldsymbol{\theta}_i)) - \phi(h_p^*(\boldsymbol{\theta}_i), \psi_{ip}) = w_{iq} h_q^*(\boldsymbol{\theta}_i) - T(w_{iq} h_q^*(\boldsymbol{\theta}_i)) - \phi(h_q^*(\boldsymbol{\theta}_i), \psi_{iq}).$$

Thus the change in social wealth created by her switching professions is

$$\begin{aligned} & (1 + e_{ip}) w_{ip} h_p^*(\boldsymbol{\theta}_i) - \phi(h_p^*(\boldsymbol{\theta}_i), \psi_{ip}) - [(1 + e_{iq}) w_{iq} h_q^*(\boldsymbol{\theta}_i) - \phi(h_q^*(\boldsymbol{\theta}_i), \psi_{iq})] = \\ & e_{ip} w_{ip} h_p^*(\boldsymbol{\theta}_i) + T(w_{ip} h_p^*(\boldsymbol{\theta}_i)) - e_{iq} w_{iq} h_q^*(\boldsymbol{\theta}_i) - T(w_{iq} h_q^*(\boldsymbol{\theta}_i)) \equiv \Delta T(\boldsymbol{\theta}_i) + \Delta E(\boldsymbol{\theta}_i), \end{aligned}$$

that is the change in the sum of her tax payments and externalities. This is exactly the discrete, career-switching analog of the intensive margin change in hours. The total change in social wealth from an increase in the marginal tax rate at y is $E[\Delta T(\boldsymbol{\theta}) + \Delta E(\boldsymbol{\theta}) | \partial\Theta(y)] f_S(y)$.

Socially optimal taxation calls for equating the sum of the intensive and career-switching effects to 0:

Proposition 1. *Optimal taxation requires that for all $y : f_S(y) > 0$, $E[\epsilon_{p^*}^h | \Theta(y)] > 0$,*

$$\frac{([E[e_{p^*} | \Theta(y)] + T'(y)] E[\epsilon_{p^*}^h | \Theta(y)] + \text{Cov}(e_{p^*}, \epsilon_{p^*}^h | \Theta(y))) f(y)}{1 - T'(y)} + E[\Delta T + \Delta E | \partial\Theta(y)] f_S(y) = 0.$$

While this intuitive result applies very generally, it provides relatively little guidance on how to map information about distribution of incomes and externality shares to optimal tax rates in the absence of detailed information on which individuals are likeliest to switch professions. To clarify the analysis further, we now consider additional assumptions that may be added to yield an especially simple formula.

⁹Note we can ignore individuals who are triply indifferent between professions as they are of measure zero even within the career switchers.

3.2 A simple case

The first such assumption is that there is no correlation between propensity to switch careers, conditional on an income level, and the externalities generated by an individual earning that income nor between the propensity to switch into a career and the externalities, conditional on the income level switched into.

Assumption 1. *For an individual in a switching set, $\partial\Theta(y)$ for some y , let $p(\theta)$ be the lower-paying of the two professions she is indifferent between and $q(\theta)$ be the higher-paying of the two professions she is indifferent between. For all y, y'*

$$E[e_{p^*}|\Theta(y)] = E[e_p|\partial\Theta(y'), w_p h_p^* = y] = E[e_q|\partial\Theta(y'), w_q h_q^* = y].$$

That is, the average individual considering (in response to a marginal change in the tax rate at y') switching either down to a lower-paying profession or up to a higher-paying profession but currently earning income y on average generates the same externality as the average individual earning that income.

The second assumption is analogous, but for the intensive margin: elasticities are uncorrelated with externalities, conditional on income.

Assumption 2. *For every y ,*

$$\text{Corr}(e_{p^*}(\theta), \epsilon_{p^*}^h(\theta)|\Theta(y)) = 0.$$

Under these assumptions, optimal taxation takes a very simple form in both the case of pure career-switching ($\epsilon_{p^*}^h(\theta) \equiv 0 \forall \theta \in \Theta$) and in the case of pure intensive reactions and no career switching ($f_S(y) = 0 \forall y$).

Proposition 2. *Suppose that, for every tax policy T , Assumption 1 holds and that all intensive elasticities are 0. Then*

$$T^*(y) = -E[e_{p^*}|\Theta(y; T)]y + T_0,$$

where T_0 is a constant across income levels. That is, up to a lump sum transfer, average tax rates are set at each income to offset the average externality created by individuals earning that income level. This is true if the average externalities are defined at the current equilibrium or at the optimal policy, as the average externalities at each income level remain the same in this case.

Alternatively suppose that, for every tax policy, Assumption 2 holds and that there is no career switching. Then

$$T^{*'}(y) = -E[e_{p^*}|\Theta(y; T^*)].$$

That is, marginal tax rates at each income level are set to offset the average externality created by individuals earning that income level given the optimal tax policy.

Intuitively, when there is pure career switching the average externality of each profession must offset the average cost of each profession because when switching across professions it is the average income rather than the marginal earnings that are relevant. On the other hand, when making marginal decisions about work, it is marginal tax rates that are relevant and thus marginal tax rates should be equated to the average externality at a given income level. This intuition is formalized in the proof of this result in Appendix B. As we discuss below, the distinction between policies generated by these two regimes is fairly small in many contexts.

Proof. See Appendix B. □

We refer to the optimal policy under pure career-switching as “average tax externality matching” (ATEM) and the optimal policy under pure intensive hours choice “marginal tax externality matching” (MTEM). To see the distinction between these policies, note that under ATEM

$$T^{\star'} = -E[e_{p^*}|\Theta(y; T^{\star})] - y \frac{\partial E[e_{p^*}|\Theta(y; T)]}{\partial y}. \quad (4)$$

Suppose that externalities are becoming larger in absolute magnitude as income rises. Then marginal tax rates are more extreme under ATEM than MTEM: if average income-conditional externalities are negative and increasing in size with income and thus marginal tax rates are positive under MTEM they will be even larger under ATEM. If average income-conditional externalities are positive and increasing in with income and thus marginal tax rates are negative under MTEM they will be even more negative under ATEM.

Proposition 2 dramatically simplifies the data requirements for determining optimal tax policy. However, it still presents two challenges. First, in the case when there are both career-switching and intensive elasticities, optimal policy is a mix of these two extremes, the mix depending on details of how large the two elasticities are over in the distribution of income. Second, with no career switching, it is the average income-conditional externality *at the optimal policy* rather than at the current equilibrium that is relevant to determine marginal tax rates. This is harder to observe from available data for obvious reasons, though this same challenge appears, and is treated as we do below, in much standard optimal tax work (such as Saez (2001)).

To confront these issues we take two approaches in calibrating our model. First, we focus primarily on the case of ATEM because we believe that career-switching elasticities are much larger than hours elasticities, at least for the most talented individuals. For example while Saez et al. (2012) argue that elasticities of intensive margin labor supply are very low (0-1 for high incomes), Goldin et al.’s evidence suggests career-switching elasticities are much higher. The share of male Harvard alumni who pursued a career in finance, for example, more than tripled from 5% in the 1969-72 cohort to 15.7% in 1989-92 cohort. Suppose that there have been no changes in inherent preferences for different careers (no labor supply shift) and thus that all of this change arises from

shifts in relative wages. Philippon and Reshef (2012) argues that wages in the financial sector have increased, over from prior to the 1980's to the late 1990's, by somewhere between 50% and 400% depending on how one adjusts for education. Average tax rates for the well-off have fallen during this period, somewhat inflating these numbers. While the Goldin et al. data only includes earnings in 2005 (rather than at comparable career stages), these are about 2.5 times as high pre-tax and likely a bit over 3 times as high post-tax in the last cohort compared to the first. Thus elasticities are somewhere between .8 and 6, between nearly a full order of magnitude and one and a half orders of magnitude greater than the intensive elasticities discussed above. Given that these increased wages were likely not fully anticipated and mostly accrued in the 1990's and 2000's and that finance's share of graduates appears, anecdotally, to have greatly increased further since the 89-92 cohort, these seem likely to be underestimates if anything. Even Keane and Rogerson (2012), who argue for higher hours elasticities, argues that longer-term elasticities (along dimensions like career choice), are likely to be much larger than intensive hour elasticities. Second, we show that our results are qualitatively robust to calibrating MTEM.

3.3 Calibration

As discussed in Section 2 above, the crucial input, in addition to distributions of income within professions and individuals across professions, to determine ATEM and MTEM policies is the externality shares of different professions. Given these, ATEM may be computed simply by, for each income, constructing the average externality using the income-conditional distribution of individuals across professions and setting the average tax rate to match this. Marginal tax rates may then be derived from average rates using equation (4). To correctly compute MTEM would require elasticities of labor supply at each income because rises in marginal tax rates will compress downwards and falls in marginal tax rates will stretch upwards the incomes earned by individuals with various average externalities. In a future draft we will explore the sensitivity of MTEM tax schedules to various assumptions about this, but given that they leave the structure of rates fixed and only change which incomes they kick in at, we here simply ignore the changes in income and match marginal tax rates to income-conditional average externalities exactly as with average tax rates under ATEM.

We do not believe compelling data exist that allow the size of externality shares to be estimated in a way likely to lead to broad agreement about these. We therefore treated these as subjective estimates likely to vary widely across individuals. We therefore did not seek, as with our other data and assumptions, to find a single, "approximately correct" assumption. Rather, we calibrated externality shares in two ways. First, we have built a flexible tool, available online at <http://home.uchicago.edu/weyl/taxapplet.html>, that allows a user to input externality shares for each occupation category and obtain a graph of MTEM and ATEM tax rates. Second, to avoid the

	Tea Party	Occupy Wall Street	Lockwood	Nathanson	Weyl
Academia/research	-.3	.6	2	.5	2
Advertising/marketing/sales	-.2	-.9	-.3	0	-.3
Agriculture	0	0	0	0	0
Arts/entertainment	-.5	.8	0	0	.1
Business operations	.3	-.2	.1	0	.1
Engineering/technical	0	.1	.4	0	4
Entrepreneurship	1.5	-.5	2	.5	2
Financial Services	.2	-.8	-.5	-.1	-.7
Law	-.8	.3	-.5	-.3	-.6
Management	.3	-.6	-.2	-.2	0
Medicine	.1	.2	.1	0	.1
Public service/military	-.8	1.5	.4	.1	.7
Real Estate	0	-.3	-.1	0	-.1
Teaching	.3	2	4	.2	4

Table 1: Externality shares in each of five calibrations.

reader having to fiddle with the tool to obtain results, we provide five calibrations that seem to us to approximate views held by two groups among the public and represent the views of the three authors of this paper. These are represented in Table 1.

The first position we consider is meant to approximate the views of a member of the Tea Party movement. The numbers assigned reflect the enthusiasm in this movement for entrepreneurship and private enterprise, skepticism of the contribution of cultural and intellectual elites and hostility to legal and and public services. The second position is meant to represent an opposite extreme, that of an Occupy Wall Street protester.¹⁰ This represents fierce hostility to financial services and other aspects of private markets typically denigrated by the left, sympathy for legal and cultural elites' value and enthusiasm for government action and education.

The next three positions represent the views of the three authors of this paper. The views of Lockwood and Weyl are essentially identical: Weyl is neutral rather than slightly negative on management, a bit more negative on law and finance, a bit more positive on public service and is slightly positive rather than neutral on the arts and entertainment. None of these make quantitatively significant differences in any results and thus we neglect Weyl's views in the rest of the paper and report only results for Lockwood's views. These views arise from taking very seriously, and perhaps to a bit of an extreme, arguments made commonly within the economics community for externalities from various professions. These include the enthusiasm for entrepreneurship and engineering and skepticism of law pushed by Murphy et al. (1991), the emphasis on the importance of research central to the modern growth literature (Acemoglu, 2008), evidence that teachers are

¹⁰It is, of course, not obvious how members of these movements would interpret their own views or whether we are capable of interpreting their views accurately. We have done our best, from anecdotal encounters (both of us know a reasonable number of members of each movement), to guess their views.

dramatically under-compensated (Chetty et al., 2011), concerns about the social value of advertising (Galbraith, 1958), a widespread recent concern that much activity in the financial sector is zero sum racing over information (Hirshleifer, 1971) or may actually be value destroying (Brunnermeier et al., 2012; Posner and Weyl, 2013b) and, most importantly, worries about managers devoting significant time to defrauding their shareholders rather than creating value for them (Shleifer and Vishny, 1997) and shareholders failing to correspondingly reduce pay due to competitive pressures (Bénabou and Tirole, 2012).

Nathanson’s position reflects largely a general Coasian skepticism of the size of externalities that persist, based on the notion that, while some externalities may persist, they most are largely internalized already by various social, public or civic mechanisms. For example, prestige that is costly either because it is limited in supply or requires scarce attention from others in society, government subsidies on various sectors and regulations on others and voluntary tithing and selective loan repayment to universities conditional on income may partly or fully offset externalities. Thus apparent lack of correspondence between material rewards of different professions, in this view, masks other rivalrous non-pecuniary or unmeasured pecuniary goods being given to different professions that internalize most externalities.

ATEM and MTEM policies under these various externality shares calculated from the data as discussed above are shown in Figure 2. The first thing to note is that ATEM yields qualitatively and often quantitatively similar results to MTEM except that 1) ATEM tends to be more extreme in both directions than is MTEM for the reasons discussed in the at the end of the previous subsection and 2) ATEM “leads” MTEM in the sense that if MTEM rates are rising (falling) in income ATEM rises first. Recall that under our assumptions MTEM is also (up to an additive constant) equal to the average tax rates under ATEM. The other properties depend critically on the assumptions about externality shares employed, and thus the ATEM and MTEM, while looking similar given an externality profile, look radically different under the various profiles discussed above.

First, consider the Tea Party shares. There are hardly any individuals with income under \$10k, and they are not of much interest for our purpose so let us focus on incomes above this level. Under both ATEM and MTEM marginal rates fall steadily from 0 to approximately -20% where they level off. In other words, marginal rates should be approximately 0 for those earning below about \$250 and roughly those above that level should be subsidized 20% on marginal income. This is driven by the fact that we assume Tea Partiers perceive positive externalities from business and finance which account for most of upper incomes and that their presence at lower incomes is offset by the presence of “parasitical” public employees and lawyers.

Second, consider the Occupy Wall Street shares. These call, under both ATEM and MTEM, for slight subsidies on those earning lower-middle incomes (\$20-40k) and then steeply rising marginal rates leveling out at just under 50% at incomes somewhere between (\$60 and \$100k) and persisting or even growing up through the full income range. This second finding is driven overwhelmingly by

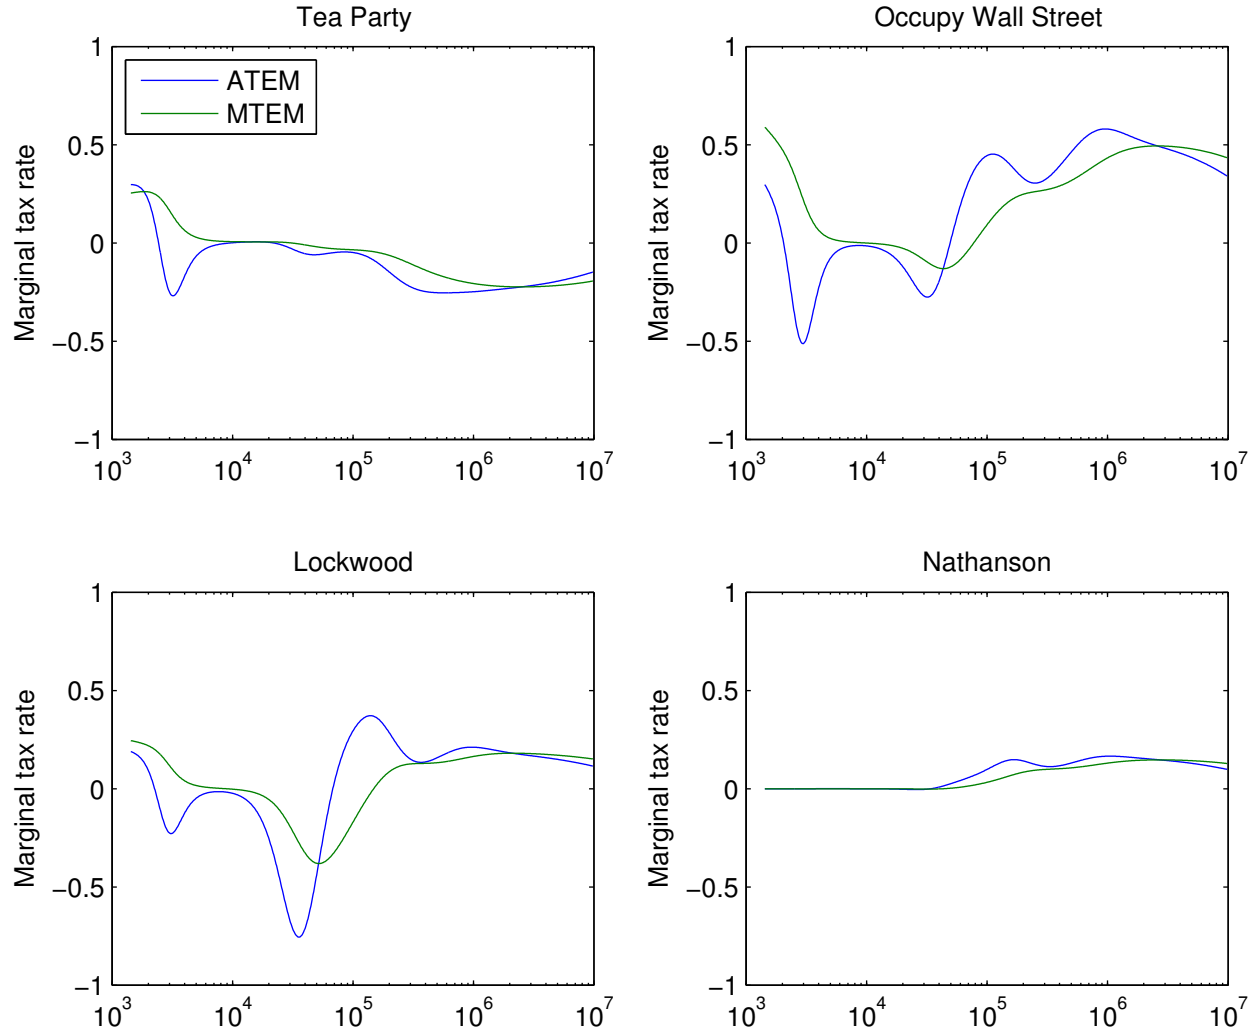


Figure 2: ATEM and MTEM marginal tax rates as a function of income plotted on a logarithmic scale for four of our five sample externality profiles.

the negative externalities we assume Occupiers perceive managers and business operators having, as well as to a smaller extent by the harms they perceive from financial services. The subsidies at middle-low incomes are largely driven by the positive externalities they perceive for teachers and public servants, many of whom have middle-low incomes.

Third, consider Lockwood's shares. In this case the qualitative patterns are similar between ATEM and MTEM, but ATEM is almost twice as extreme and because it is more striking we focus on this case. The striking pattern here is a dip of marginal rates to very low levels and then their rise to quite high levels before they level out around 20%. The dip begins around \$20k at a rate near 0, reaches its nadir at around \$35k and a range of around -75%. It then rises, crossing 0 at an income close to \$70k, peaks at an income of about \$150k and a rate of 40%. At this point it begins to fall and levels out at around 20% at an income of about \$250k. The Lockwood pattern is driven by the

fact that only at middle incomes do the professions (engineering/technical, public service/military, academia/research) that Lockwood believes have positive externalities have a strong representation, making subsidies to this group attractive. Above this middle income range, financiers and especially managers come to dominate, which Lockwood believes have a small negative externality.

Finally, the results for Nathanson’s shares are the easiest to explain. He believes few externalities are significantly different from 0 and obtains under both ATEM and MTEM marginal rates never significantly above 10%, and never negative. A very small marginal rate on the wealthy is driven by the small negative externality from management and finance he perceives.

While simple and transparent, the assumption of no correlation between externalities and tendency to switch into and out of a career may significantly bias optimal marginal tax rates. The reason is the strong correlation between income and externality share under several of the externality share profiles. It seems likely that when individuals switch professions they tend to stay in a similar quantile of the distribution of income in that profession, because their ability level is relatively portable. Thus it is highly unlikely that raising marginal rates at \$300k would lead, say, an academic or public servant earning an exceptionally high salary of, say, \$400k to switch careers; it is much more likely it would lead a financier earning the rather modest-for-finance wage of \$400k to switch down to a lower wage in academia or public service.

4 Discussion

For clarity and brevity above we tried to develop our theory with limited reference to external concepts. However, our theory is closely related to, and has implications for, several other literatures that we now discuss.

4.1 Labor supply elasticity debate

An important recent controversy in public finance concerns the elasticity of labor supply. One literature, surveyed by Saez et al. (2012), highlights that the short-term elasticities of labor supply and taxable income that can be measured through natural experiments are low (on the order of .1 – .5 for high income earners) and largely driven by welfare-irrelevant evasion or inter-temporal substitution. A second literature, surveyed by Keane (2011) and Keane and Rogerson (2012), argues that this evidence is consistent with large long-term elasticities of labor supply that would not appear in short-term estimates and that such large long-term elasticities also help rationalize international data. The first literature tends to use low elasticities to minimize the deadweight loss from taxation and argue for highly redistributive tax policy; the second literature argues that taxes are highly distortionary.

Our work offers a third alternative. On the one hand, our calibration of career-switching elastic-

ities is largely consistent with Keane and Rogerson arguments for high (perhaps even higher than they claim) long-term labor-supply elasticities among talented and wealthy individuals. On the other, it suggests that high long-term elasticities (interpreted as career-switching elasticities) need not imply lower optimal marginal tax rates. In fact, they may make optimal marginal rates higher over many ranges if high marginal rates are implemented for efficiency rather than equity reasons.

Thus it emphasizes that the externalities of various professions, rather than the elasticities of labor supply, may be the primary determinants of optimal income taxation. If externalities are small or more positive at higher incomes, given our calibrated career-switching elasticities, optimal tax rates are likely to be low even with a significant redistributive motive as we discuss further in Subsection 5.3. On the other hand, if externalities are large and more negative at higher incomes, optimal tax rates on the wealthy are likely to be very high, increasing in the long-term elasticity and not very sensitive to redistributive motives.

Given this, the relative empirical attention these two sets of parameters have received seems disproportionately weighted towards labor supply elasticities. Thousands of paper, a few hundred of which are surveyed in the papers discussed above, have been devoted to measuring the elasticity of labor supply and of taxable income more broadly. On the other hand, we are not aware of a single paper that aims to measure the externalities created on average by professions such as law, engineering, entrepreneurship, research or teaching. Even those that are indirectly informative about these topics are limited; some of these were discussed in the Subsection 3.3 above. Our theory suggests this topic deserves more attention in future research.

4.2 Agency and multi-tasking

While the optimality of softening incentives to promote efficient allocation of labor is, to our knowledge, largely unexplored in optimal tax theory, it is widely understood in corporate finance and agency theory.¹¹ In particular, Holmström and Milgrom (1991) argue that strong material incentives for one, observable dimension of work effort may reduce effort along unobservable dimensions to the extent that effort along the two dimensions is substitutable. Our theory is similar, except that the substitutability arises from the (un-modeled and empirically calibrated) correlation between income levels and the externalities of professions. This alternative micro-foundation of their model also acts as a basis for calibrating it empirically in a taxation context; we are not aware of any analogous calibration of the optimal contract schedule in the original agency context in which they proposed their model.¹² If the productivity, the private payoff from and the allocation of time to

¹¹A notable exception is the work of Piketty et al. (2011), who argue that taxation may reduce effort expended bargaining for higher compensation. While we agree with the spirit of this result and are motivated by the macroeconomic evidence the authors present, we believe that compensation bargaining has a small elasticity compared to career choice and a career-based theory offers a more useful basis for calibration.

¹²Slade (1996) tests directional comparative statics of the Holmström and Milgrom model, but does not structurally calibrate optimal policy. As far as we know the only other paper to exploit the equivalence of agency and tax theory

tasks them could be identified within firms, the theory could be applied with data on the allocation of time to these tasks.¹³

4.3 Allocation of talent

Baumol (1990) and Murphy et al. (1991) emphasize the importance of the allocation of talent for long-run growth and claim that externalities in some of the professions we highlight are similar to the views of Lockwood and Weyl. While a considerable literature builds on this analysis, little work has discussed policy tools that might be used to improve the allocation of talent. Acemoglu (1995) discusses long-term cultural factors and the possibility of multiple equilibria and path dependency, but does not highlight policy tools that may be used to shift equilibria.

More closely related to our work is that of Philippon (2010) and Rothschild and Scheuer (2012).¹⁴ Both papers investigate public policies aimed at reallocating talent. Philippon considers the use of taxation to affect the allocation of talent between a financial and entrepreneurial sector in a model where financiers serve as conduits for funding of research and entrepreneurship. By contrast, externalities are directly assumed (rather than arising endogenously) in our framework and we consider many professions simultaneously. Furthermore Philippon focuses on sector-specific instruments rather than horizontally equitable policies.

Rothschild and Scheuer are closer to our work in this dimension, considering horizontally equitable optimal tax policy. Our two papers focus on aspects and approaches to this problem. They restrict attention to two professions (one with negative externalities and one with none), but allow for richer targeting of externalities: in their model, rather than externalities accruing uniformly across the population they may be targeted either at individuals within the “rent-seeking” sector or towards individuals in the productive sector.¹⁵ Like Philippon they emphasize the, perhaps counter-intuitive theoretical finding that it may be theoretically optimal for policy to (implicitly or explicitly) subsidize the unproductive wealthy, in their case because the negative externalities within the unproductive sector discourages further wasteful entry into this sector.

By contrast our emphasis is quantitative: our goal is to determine how the magnitude of optimal taxes varies with assumptions about externalities, allowing for the many professions that make up real-world top income distributions and for positive as well as negative externalities.¹⁶ Given the

to link of quantitative optimal tax work to the largely theoretical agency literature is Prendergast (2012), albeit in the more classical context of the Mirrleesian insurance-incentives trade-off.

¹³For example, a friend of one of the authors who works at an investment bank reports that, “I spend one-third of my time creating profits for the firm, one-third of my time ensuring I get credit for those profits and one-third of my time ensuring that I get paid for the profits I got credit for.” Whatever the actual proportions, presumably optimal compensation structure is highly sensitive to these proportions.

¹⁴The original version of the paper that this draft was based on, “Psychic Income, Taxation and the Allocation of Talent”, was written prior to a draft of either of those papers. However, none of those authors cites or otherwise appears to be aware of that work and we thus assume their work was entirely independent.

¹⁵Rothschild and Scheuer also include a redistributive motive, which we hope to do in a future draft.

¹⁶As a result of these contrasting goals, we take a very different methodological approach to Rothschild and Scheuer

lack of empirical information we have on the targeting of externalities, this requires us, among other restrictions, to either ignore the effects highlighted by Philippon, Rothschild and Scheuer or to treat targeting as another subjective input. We chose the former course because we believe (and suspect many readers will as well) that the negative externalities in a sector are just as likely to raise as to lower the marginal product from others entering that sector (more high speed traders require other investors to make greater investments to avoid being front-run, thereby hiring more financiers and the legal profession hosts similar arms races). The dramatic rise of the financial sector the United States seem over the last thirty years as documented by Philippon and Reshef (2012), Goldin et al. (2013) and Philippon (2012) seems to belie the equilibrating forces Philippon (2010), Rothschild and Scheuer emphasize. Thus we adopt what Rothschild and Scheuer call the “naïvely Pigouvian” perspective that wages, externalities and entry into professions are independent. An interesting extension of our work would be to incorporate the targeting emphasized by their work into a quantitative framework such as ours

To get a rough sense of the importance of the allocation of talent that our taxes aim to address, consider the Goldin et al. data. Using the Lockwood externality shares, we computed the (2005) income-weighted average externality share in a subset of the Goldin data (about 70-80% depending on cohorts) who pursued careers in one of academia/research, financial services, law, medicine or management for the 1969-72 and 1989-92 cohorts. This average externality share fell from -.18 to -.34. Suppose that this shift took place for the full top 5% of the income distribution, who capture about 35% of income according to Atkinson et al. (2011).¹⁷ This shift would then imply a $.17 \cdot .35 \cdot 100 \approx 6\%$ of GDP shift in national income in some combination of reducing aggregate production increasing and the private returns of those at the top. If all were the latter (private returns rose and aggregate production stayed constant) this would account for nearly all of the increase in the share of national income appropriated by the top 5% between 1969 and 1989 and nearly half of the increase between the trough of this share in the early 1970’s and its peak in the mid-2000’s. If all were reduction in social product, it would account for around a .25% reduction in GDP growth each year for 25 years, roughly one sixth of the approximate 1.5 percentage point reduction in US GDP growth between between the “Golden Age” (1949-1973) and the “Washington Consensus” (1982-2007) periods.

as well. For example, we use the taxation principle in the spirit of Saez (2001) rather than the revelation principle of Mirrlees (1971). Similarly we make assumptions (about substitution patterns across professions) that are convenient for calibration to the data we have available, rather than for illustrating possibilities.

¹⁷It seems likely that that the shift among Harvard alumni was more extreme than in the rest of the population. However, it also seems likely that the shift among these students was greater by the mid-2000’s than it was by 1989-92, and greater from the mid-1950’s than it was since the late 60’s. We suspect these two effects close to balance.

4.4 Debates on taxation outside neoclassical economics

Public debate over tax policy rarely focuses on the parameters emphasized in optimal tax theory (viz. the degree of inequality or the responsiveness of work to taxes). Instead, as Mankiw (2010) notes, much rhetoric instead focuses on whether the rich “deserve” the wealth they have accumulated. The left attempts to delegitimize the wealth of the rich (claiming it is misbegotten or crooked) and of the right to hold the wealthy up as job creators and entrepreneurs. Some simple public opinion data (Parker, 2012) is suggestive here. While 55% of Republicans believed the rich were more likely than others to be hardworking and 18% believed they were more likely to be honest, only 33% and 8% of Democrats agreed on each count respectively. On the other hand only 42% of Republicans believe the rich are more likely than others to be greedy, while 65% of Democrats do. For a future draft of this paper we hope to collect text and/or further public opinion data to confirm this more rigorously. However, it seems clear that much public debate focuses on the degree of externalities created by the professions in which the rich are employed rather than on elasticities.

None of this is new: in the nineteenth century reformers accused the wealthy of being “robber barons” and Marx (1867) accused the wealthy not of being unresponsive to tax rates but of being unproductive exploiters of the truly productive working class. Literature on the right, such as Rand (1957), largely consists of hagiographic representations of the social contributions of the wealthy rather than depictions of their willingness to shirk if taxes rise. In fact that book is largely devoted to the unwillingness of the rich to shirk even when they are nearly enslaved.

Our theory provides a natural formal, quantitative language for these debates which fit only unnaturally with the Mirrleesian framework. In so doing, data can be brought more easily to bear; disputes become questions of which professions exactly are claimed by different sides to have different degrees of externalities and which professions in fact earn which incomes. Such questions should be easier to settle empirically than are broad and vague claims about the moral worth of different social groups. At very least it suggests that if economics wants to speak to these public debates it should focus more attention on the extent and degree of these externalities.

5 Extensions

In Section 3 we developed sharp, simple and intuitive characterization of optimal implicit Pigouvian taxation of income under a stark set of assumptions. In this section we discuss the robustness of the results to adding several realistic features to the model using a combination of analytical and computational techniques. While our characterization is less clean in these cases, our qualitative results remain and are in many cases strengthened by these adaptations.

5.1 General ability: analytic results

The model in Subsections 3.2 and 3.3 assumes a strong form of orthogonality between elasticities and externalities conditional on income level. One natural scenario in which this would be contradicted, along the career-switching margin, is that there exists a general ability possessed by each individual which determines her wage in all professions. In this case, individuals are likely to systematically shift from professions with high average earnings to those with low average earnings when taxes rise. An academic making a million dollars a year could likely earn far more were she to switch professions; she is just exceptionally talented. It is thus almost inconceivable that when marginal tax rates change she will switch to being a poorly compensated financier. On the other hand a financier earning a million dollars a year may not be so exceptionally talent and might be willing, if marginal tax rates rose, to switch into academia at a lower salary. The reason is that both individuals likely face the same *ranking* of wages in the two professions and thus marginal tax rates are likely to cause only switches *down* from generally high-paying to generally low-paying professions, not the reverse.

The same scenario is also likely to imply differential intensive elasticities of hours supply *conditional on income*. An academic making a million dollars a year is likely to be working many more effective hours than is an financier earning that salary both because her wage in finance is likely to be lower (and thus, to be earning the same income, she must be working harder) and because her non-pecuniary cost of work is likely to be lower. She is therefore likely to be pressing herself to the limit of her exertions and thus to be much less elastic to an increase in her monetary compensation than is the financier. Thus one should expect individuals in a high-paying profession to be more elastic, *conditional on income*, to compensation than are those earning the same income in a generally low-paying profession.

While few simple analytical results may be obtained in general when we depart from this story, under a stark simplification of this story natural result is possible. In particular, first suppose that there is no intensive elasticity, that ability is fully portable across careers in the sense that individuals switch between the same quantile of the distribution of profession-conditional income when they switch, that there are exactly two profession and that the profession-contingent income in these two are strictly ranked by first-order stochastic dominance. Then marginal tax rates will always be more sensitive to externalities than those given by ATEM in the sense that, starting at ATEM tax rates, there is a first-order welfare gain from raising marginal tax rates at every point if the high-paying profession has more negative externalities and a first-order gain from lowering them if the high-paying profession has more positive externalities. Second, suppose that there is no career-switching, that the curvature of effort cost is only a function of hours and is an increasing function of those, that everyone in the lower paid profession has a higher non-pecuniary benefit of working (in a first-order stochastic dominance sense) and that the marginal cost of working decreases in the non-pecuniary benefit of work. Then marginal tax rates will always be more sensitive to externalities than those given by MTEM in the same sense.

More formally, let us first return to the basic model, without intensive elasticities, and suppose that $n = 2$ and that in profession p there is a reference distribution of income with CDF $G_p(w)$ and PDF $g_p(w)$. Given that there are no intensive elasticities, we simply assume every individual works one hour and thus her wage in a profession is equal to her income in that profession. Every individual i is endowed with a uniformly distributed general ability $a_i \in [0, 1]$. They also receive a vector of non-pecuniary total costs of work in each profession ϕ_i forming a type $\theta_i = (a_i, \phi_i)$. Individual i earns wage $G_p^{-1}(a_i)$ in profession p . Otherwise the model is exactly as in Subsection 3.1.

Proposition 3. *Suppose that G_1 first-order stochastically dominates G_2 . If $e_1 < e_2$ then beginning from ATEM, there is a first-order welfare gain from raising the marginal tax rate at each wage level w in the support of the wage distribution for which $\partial\Theta(w)$ includes individuals switching either from or to a wage level including individuals from both professions. On the other hand, if $e_1 > e_2$ then beginning from ATEM, there is a strictly positive first-order welfare gain from lowering the marginal tax rate at each wage level w in the support of the wage distribution for which $\partial\Theta(w)$ includes individuals switching either from or to a wage level including individuals from both professions.*

Proof. Suppose that $T(w) = T_0 - E[e_{p^*}|\Theta(w)]w$. Consider any w meeting the description in the proposition statement. Partition $\partial\Theta(w)$ into two sets: $\widetilde{\partial\Theta}(w)$ and $\overline{\partial\Theta}(w)$ where $\widetilde{\partial\Theta}(w)$ includes only types switching between wage levels at least one of which includes a positive density of individuals of each profession and $\overline{\partial\Theta}(w)$ includes only individuals switching between wage levels occupied each by a sole profession. By assumption, $\widetilde{\partial\Theta}(w)$ is non-empty; let

$$\lambda(w) \equiv \frac{\int_{\widetilde{\partial\Theta}(w)} f(\theta) d\theta}{\int_{\widetilde{\partial\Theta}(w)} f(\theta) d\theta + \int_{\overline{\partial\Theta}(w)} f(\theta) d\theta}.$$

Then

$$\begin{aligned} E[\Delta T + \Delta E|\partial\Theta(w)] &= \lambda(w) E\left[T(w_{ip}) + e_{ip}w_{ip} - T(w_{iq}) - e_{iq}w_{iq}|\widetilde{\partial\Theta}(w)\right] + \\ &\quad [1 - \lambda(w)] E\left[T(w_{ip}) + e_{ip}w_{ip} - T(w_{iq}) - e_{iq}w_{iq}|\overline{\partial\Theta}(w)\right]. \end{aligned}$$

Note that the second term is 0 by the argument in the proof of Proposition 2. On the first term, note the for any type in $\widetilde{\partial\Theta}(w)$ we have

$$\begin{aligned} T(w_{ip}) + e_{ip}w_{ip} - T(w_{iq}) - e_{iq}w_{iq} &= \\ -E[e_{p^*}|\Theta(w_{ip})]w_{ip} + e_{ip}w_{ip} + E[e_{q^*}|\Theta(w_{iq})]w_{iq} - e_{iq}w_{iq} &< (>)0 \end{aligned}$$

if $e_1 < (>)e_2$ as $p = 1$ and $q = 2$ always because no individual has a higher income in profession 2 than 1 by the general income assumption and first-order stochastic dominance. \square

While neither of these results directly implies marginal rates are higher at every point than

under Proposition 2, they are strongly suggestive in this direction. This sort of result also turns out to hold in our computational results below.

Next consider the case of no career-switching. That is, suppose there are two professions and every individual is stuck in one profession or the other. Let G_i be, as above, the wage distribution in profession i (but now it is not simply a reference distribution as there is no career-switching). Our proposition depends on the following assumptions.

- Assumption 3.** 1. $\phi(h; \psi) = \phi(h) - \gamma\psi h - \nu\psi$ and $\gamma, \nu > 0$. That is, the curvature of disutility from work depends only on the number of hours worked and psychic income from work makes working less painful, as well as making the job more attractive overall.
2. Conditional on income, the distribution of ψ in profession 2 first-order stochastically dominates that in profession 1. That is, individuals typically prefer working in profession 2, all else equal.¹⁸

Under these conditions, we have a proposition similar to Proposition 3 except that the baseline policy, given that there is no career-switching, is MTEM rather than ATEM.

Proposition 4. Suppose that the conditions of Assumption 3 hold. If $e_1 < e_2$ then beginning from MTEM, there is a first-order welfare gain from raising the marginal tax rate at each income level y in the support of the income distribution for which elasticities are positive for a positive measure of individuals at that wage. On the other hand, if $e_1 > e_2$ then beginning from MTEM, there is a strictly positive first-order welfare gain from lowering the marginal tax rate at each income level y in the support of the income distribution for which elasticities are positive for a positive measure of individuals at that wage.

Proof. See Appendix B. □

5.2 General ability: computational results

In this subsection we explore the robustness of the results of the previous section computationally in a richer setting but with stronger structural restrictions on primitives. In particular, given our limited data on elasticities, we here assume away intensive elasticities and impose the common logistic functional form on non-pecuniary utilities, while matching the distribution of wages in professions to the observed equilibrium income distribution non-parametrically, calibrating the remaining single parameter using the career-switching elasticity we calibrated above and maintaining our “general ability” assumption to calculate counter-factual wages for individuals in other professions. We then show, as suggested by the previous subsection, that our results are only strengthened by this change to the baseline model of Subsections 3.2 and 3.3.

¹⁸This can be derived from first-order stochastic dominance conditional on wages and an appropriate condition on the ranking of wages.

As in the pure career choice model in previous subsection, each individual is endowed with a fully-portable, general ability a_i that would entitle her to pecuniary income from that quantile of the reference income distribution of any of the professions. Each profession has a mean non-pecuniary income and each individual has, in addition to this, for each profession an independently and identically distributed (across professions and individuals) mean-zero Type I Extreme Value (Gumbel) component of her non-pecuniary income in that profession. The standard deviation of these idiosyncratic draws, β , determines how responsive individuals are to changes in the attractiveness of different professions. Intuitively, if β is large allocation across professions is primarily driven by individual idiosyncrasies and thus there will be little individual response to changes in the relative wages of different professions and conversely if β is small.

We estimate the model, we use the elasticity of individuals into the finance profession in response to changes in relative wages discussed at the end of Subsection 3.2 above, choosing the quite conservative elasticity of 1, to calibrate β based on the behavior of the financial profession given its current share; under our logit choice assumption, this elasticity is equal to $\beta(1 - s)$ where s is the share of individuals in finance. We assumed, inconsistently with what our model predicts, that the reference distribution of income in each profession is just the observed empirical distribution of income in that profession. While this is clearly incoherent and we plan to estimate these distributions semi-parametrically from our model in a future draft, we suspect the basic results are unlikely to be driven by this kludge and therefore we report results based on it as a first pass.

This model obviously has many limitations: ability is fully portable, disallowing comparative advantage; substitution patterns across professions obey the implausible independence of irrelevant alternatives assumption; there are no intensive elasticities, etc. Perhaps worst of all, exactly how these assumptions impact our results is far from transparent. A more realistic model of any of these features would require much more detailed data than we have access to. We use the model, however, not as a realistic representation of any of these features. Instead, our goal is just to provide an example model with more than two professions where under the general ability assumption we obtain a result in the spirit of that demonstrated in general in the simpler case: that optimal marginal rates are more responsive to externality share assumptions than implied by our baseline ATEM analysis.

To do this, in our computational model we searched computationally for optimal piecewise linear rates over the range \$50k to \$250k (the main range where the disagreements between our various externality share assumptions manifested in Subsection 3.3) with each piece covering \$25k of income. We restricted marginal rates to be below 100% to avoid incentives for money burning; this restriction was not generally important, except for the Occupy Wall Street preferences where rates were essentially always fully confiscatory. We thus discarded this simulation and only report those for the Tea Party, Lockwood and Nathanson.

These are pictured in Figure 3. The first thing to note is that the discreteness, coupled with




Figure 3: Simulated optimal piece-wise constant marginal tax rates for incomes between \$0 and \$250k in our computational general ability model for three externality share profiles (Tea Party on the left, Lockwood in the middle and Nathanson on the right).

the poor quality of the optimization algorithm we have thus far been able to run, leads to a lot of choppiness in the results we obtain. Nonetheless the basic message is fairly clear: the results are qualitatively similar to, but much more striking and extreme than, those we obtain in Subsection 3.3. The Tea Party externality shares (somewhat ironically) in the left-most panel of Figure 3 as above lead to negative marginal rates, but here they are much lower than before, averaging around -50% rather than -20% above and sometimes reaching much lower. Lockwood’s shares call for marginal rates above 50% on most incomes above \$100k. Even Nathanson’s highly skeptical position supports non-trivial marginal rates, around 30%, on the rich. This provides suggestive evidence that the effects developed theoretically in the previous subsection may be quantitatively important and thus that our main results from Subsection 3.3 may significantly *understate* how widely optimal marginal rates vary as a function of externality shares.

5.3 Redistribution

In a future draft of this paper, we hope in this subsection to include a version of the model from the previous subsection that includes intensive elasticity and a redistributive motive under which the marginal social welfare weight placed on any individual is inversely proportional to their income. This will allow us to introduce redistribution in the spirit of Diamond (1998) and Saez (2001) while maintaining our simple quasi-linear structure for agent preferences. First, in this subsection, we hope to show the basic robustness of our results.

5.4 Sensitivity to elasticities and externalities

Then, in this next subsection, we hope to use this to show the sense in which optimal tax rates are more sensitive to externalities than they are to elasticities of labor supply.

5.5 Quantitative welfare gains

In this subsection we plan to quantify various aspects of the welfare gains from optimal policy. First we hope to quantify the gain from moving from current policy to optimal policy under our computational model from Subsection 5.2. We then consider how much more welfare could be gained by profession-specific taxation. Finally we compare these welfare gains to those in the Saez (2001) model.

6 Conclusion

This paper proposes an alternative framework for the optimal taxation of top incomes to the redistributive theory of Mirrlees (1971). Income taxation acts as an implicit Pigouvian tax that is used to reallocate talented individuals from professions that cause negative externalities to those that cause positive externalities. Optimal tax rates are highly sensitive to which professions generate what externalities. If externalities are as large as Lockwood and Weyl believe, the worsening allocation of talent in the United States that higher taxation could mostly address is large enough to account for half of the increase in inequality or one sixth of the fall in growth between the 1945-1970 period and the 1980-2005 period.

Our results and the assumptions driving them naturally suggest several directions for future research. First, as we hope to explore in a future draft of this paper, we conjecture that a large part of the public debate over taxation is driven by views about externalities rather than the traditional Mirrlees parameters (elasticities and distributive preferences). We hope to collect data, partly through a freely available web application at <http://home.uchicago.edu/weyl/taxapplet.html> that plots optimal taxes based on views about externalities, about the views that different individuals have about these parameters.

Second, we assumed that externalities are homogeneous within a profession. However, in reality, externalities are highly heterogeneous within professions. As Mankiw and Whinston (1986) emphasize, entrepreneurial firm formation may be excessively supplied if firms are simply imitating the products of existing firms just as Hirshleifer (1971) emphasized that high-speed trading is over-supplied, while Posner and Weyl (2013a) show that long-term price discovery of large bubbles is just as likely to be undersupplied as are innovative breakthroughs. Thus many of the largest gains may come from reallocations within a profession between productive and unproductive activities.

Uniform income taxation is unlikely to be a sufficient tool to achieve this reallocation. Mechanisms that do are an exciting direction for future research.

Third, while we and others have personal views about externalities, we have no systematic strategy for measuring these empirically. Plausible strategies will almost certainly be sector specific: in finance the relevant concerns will likely be high-speed trading and risk-increasing speculation, while in law they might be the extent to which better judicial outcomes are facilitated by higher spending. On the positive side, measuring entrepreneurial externalities from the consumer surplus generated by entry will likely rely on the demand estimation typically used in industrial organization, while measuring the under-compensation of teachers will likely rely on standard techniques for measuring the returns to education. For those interested in optimal tax rates, our paper suggests that these measurements may be crucial.

Finally, we assumed that profession-specific taxation was, for a variety of reasons, infeasible. Greater exploration of the empirical relevance of these reasons and whether sector-specific taxation could be implemented in a relatively efficient manner is an important topic for future research.

References

- Acemoglu, Daron**, “Reward Structures and the Allocation of Talent,” *European Economic Review*, 1995, *39* (1), 17–33.
- , *Introduction to Modern Economic Growth*, Princeton: Princeton University Press, 2008.
- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez**, “Top Incomes in the Long Run of History,” *Journal of Economic Literature*, 2011, *49* (1), 3–71.
- Bakija, Jon, Adam Cole, and Bradley T. Hei**, “Job and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from U. S. Tax Return Data,” 2012. <http://web.williams.edu/Economics/wp/BakijaColeHeimJobsIncomeGrowthTopEarners.pdf>.
- Baumol, William J.**, “Entrepreneurship: Productive, Unproductive and Destructive,” *Journal of Political Economy*, 1990, *98* (5), 893–921.
- Bénabou, Roland and Jean Tirole**, “Bonus Culture: Competitive Pay, Screening, and Multitasking,” 2012. <http://www.princeton.edu/~rbenabou/papers/Multitasking>
- Brunnermeier, Markus K., Alp Simsek, and Wei Xiong**, “A Welfare Criterion for Models with Biased Beliefs,” 2012. http://scholar.princeton.edu/markus/files/welfare8_g.pdf.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff**, “The Long-Term Impacts of Teachers: Teacher Value-Added and Student Outcomes in Adulthood,” 2011. http://obs.rc.fas.harvard.edu/chetty/value_added.pdf.

- Diamond, Peter A.**, “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates,” *American Economic Review*, 1998, 88 (1), 83–95.
- Diamond, Peter and Emmanuel Saez**, “The Case for a Progressive Tax: From Basic Research to Policy Recommendations,” *Journal of Economic Perspectives*, 2011, 25 (4), 165–190.
- Galbraith, John Kenneth**, *The Affluent Society*, New York: Houghton Mifflin, 1958.
- Goldin, Claudia, Lawrence F. Katz, Naomi Hausman, and Bryce Ward**, “Harvard and Beyond Project,” Survey Data Project, Harvard University 2013.
- Harsanyi, John C.**, “Cardinal Utility in Welfare Economics and in the Theory of Risk-taking,” *Journal of Political Economy*, 1953, 61 (5), 434–435.
- Hirshleifer, Jack**, “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 1971, 61 (4), 561–574.
- Holmström, Bengt and Paul Milgrom**, “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics & Organization*, 1991, 7 (Special Issue), 25–52.
- Kaplan, Stephen N. and Joshua D. Rauh**, “Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes,” *Review of Financial Studies*, 2010, 23 (3), 1004–1050.
- Keane, Michael and Richard Rogerson**, “Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom,” *Journal of Economic Literature*, 2012, 50 (2), 464–476.
- Keane, Michael P.**, “Labor Supply and Taxes: A Survey,” *Journal of Economic Literature*, 2011, 49 (4), 961–1075.
- Mankiw, N. Gregory**, “Spreading the Wealth Around: Reflections Inspired by Joe the Plumber,” *Eastern Economic Journal*, 2010, 36, 285–298.
- Mankiw, N. Gregory and Michael D. Whinston**, “Free Entry and Social Inefficiency,” *The RAND Journal of Economics*, 1986, 17 (1), 48–58.
- , **Matthew Weinzierl, and Danny Yagan**, “Optimal Taxation in Theory and Practice,” *Journal of Economic Perspectives*, 2009, 23 (4), 147–174.
- Marx, Karl**, *Das Kapital, Kritik der Politischen Ökonomie*, Hamburg, Germany: Verlag von Otto Meissner, 1867.
- Mirrlees, J. A.**, “An Exploration in the Theory of Optimal Taxation,” *Review of Economic Studies*, 1971, 38 (2), 175–208.

- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny**, “The Allocation of Talent: Implications for Growth,” *Quarterly Journal of Economics*, 1991, 106 (2), 503–530.
- Parker, Kim**, “Yes, the Rich Are Different,” *Pew Social and Demographic Trends*, 2012, August 27.
- Philippon, Thomas**, “Financiers versus Engineers: Should the Financial Sector be Taxed or Subsidized?,” *American Economic Journal: Microeconomics*, 2010, 2 (3), 158–182.
- , “Has the U.S. Finance Industry Become Less Efficient? On the Theory and Measurement of Financial Intermediation,” 2012. http://pages.stern.nyu.edu/~tphilipp/papers/Finance_Efficiency.pdf.
- and **Ariell Reshef**, “Wages and Human Capital in the U.S. Finance Industry: 1909–2006,” *Quarterly Journal of Economics*, 2012, 127 (4), 1551–1609.
- Pigou, Arthur C.**, *The Economics of Welfare*, London: MacMillan, 1920.
- Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva**, “Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities,” 2011. <http://emlab.berkeley.edu/users/saez/piketty-saez-stantchevaNBER11thirdelasticity.pdf>.
- Posner, Eric and E. Glen Weyl**, “Benefit-Cost Analysis for Financial Regulation,” *American Economic Review Papers and Proceedings*, 2013, 103 (3).
- and —, “An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to 21st Century Financial Markets,” *Northwestern University Law Review*, 2013, 107 (3).
- and —, “Pooled Legal Finance: Restoring Efficiency and Equity to the Legal System,” *University of Chicago Law Review*, Forthcoming.
- Prendergast, Canice**, “The Empirical Content of Pay-for-Performance,” 2012. Working paper, University of Chicago Booth School of Business.
- Rand, Ayn**, *Atlas Shrugged*, New York: Random House, 1957.
- Rawls, John**, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971.
- Rothschild, Casey and Florian Scheuer**, “Optimal Taxation with Rent-Seeking,” 2012. http://www.stanford.edu/~scheuer/rent_seeking.pdf.
- Saez, Emmanuel**, “Using Elasticities to Derive Optimal Income Tax Rates,” *The Review of Economic Studies*, 2001, 68 (1), 205–229.

- , **Joel Slemrod**, and **Seth H. Giertz**, “The Elasticity of Taxable Income With Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Literature*, 2012, 50 (1), 3–50.
- Shleifer, Andrei and Robert W. Vishny**, “A Survey of Corporate Governance,” *A Survey of Corporate Governance*, 1997, 52 (2), 737–783.
- Slade, Margaret E.**, “Multitask Agency and Contract Choice: An Empirical Exploration,” *International Economic Review*, 1996, 37 (2), 465–486.
- Veiga, André and E. Glen Weyl**, “Multidimensional Product Design,” 2012. http://home.uchicago.edu/weyl/Product_design_5_12.pdf.
- Wilson, Robert B.**, *Nonlinear Pricing*, Oxford: Oxford University Press, 1993.

Appendix

A Data Construction

Classification

We adopt the job classifications used in Bakija, Cole, and Heim (2012) (henceforth BCH) with a few exceptions. They use both SOC and NAICS classifications to define jobs. Their classification scheme is in Table A.1 of their paper. Our changes to this scheme are the following:

- We do not exclude real estate and construction NAICS when constructing ”skilled sales”;
- We use all SOC codes of the form 27-xxxx for ”Arts, media, and sports”;
- We exclude primary/secondary teachers from ”Government, teachers, social services” and define it as a separate category as all SOC codes between 25-2000 and 25-3999;
- We re-define ”Medical” as all SOC codes in the set $[29-1011, 29-1029] \cup \{29-1041\} \cup [29-1061, 29-1069] \cup \{29-1081\} \cup \{29-1131\}$.

Income Distribution Estimation

The BLS earnings data provide the 10-25-50-75-90 income quantiles for each SOC-NAICS pairing. We estimate a lognormal distribution for each SOC-NAICS pairing using these quantiles in the following way. Let μ and σ be the mean and standard deviation of the corresponding normal

distribution. Let y_j be the given percentiles, for $j = 0.1, 0.25, 0.5, 0.75, 0.9$. We use $\hat{\mu} = e^{y_{0.5}}$ and then use

$$\hat{\sigma} = \arg \min \sum_j (F(y_j; \hat{\mu}, \hat{\sigma}) - j)^2$$

where $F(; \mu, \sigma)$ is the CDF of a lognormal with parameters μ and σ . If we lack the data on $y_{0.5}$, we use just $y_{0.1}$ and $y_{0.25}$, for which there is a unique solution for μ and σ , and if we lack $y_{0.25}$ as well, we drop the observation.

Let i index our job classifications and let k index the SOC-NAICS pairings within a job classification. Let n_k be the number of workers in a SOC-NAICS pairing, which is provided by the BLS data. We construct a naive estimate of the income distribution CDF F_i in job classification i with

$$\hat{F}_i = \frac{\sum_k n_k F(; \hat{\mu}_k, \hat{\sigma}_k)}{\sum_k n_k}.$$

This naive estimate essentially estimates the lower portion of the true CDF F_i . We estimate the full CDF using the following data from BCH. For two income levels y_h and y_{hh} , corresponding to the top 1% and top 0.1% of the income distribution for the United States, BCH provide the shares $s_{i,h}$ and $s_{i,hh}$ of workers in industry i with income exceeding \$295,000 and \$1,246,000 who work in job i . Combining these with the quantiles generated from BLS data yields a vector of seven quantiles, $j_i = (0.1, 0.25, 0.5, 0.75, 0.9, 1 - s_{i,h}, 1 - s_{i,hh})'$ and corresponding income thresholds,

$$y_i = \begin{pmatrix} \hat{F}_i^{-1}(0.1) \\ \hat{F}_i^{-1}(0.25) \\ \hat{F}_i^{-1}(0.5) \\ \hat{F}_i^{-1}(0.75) \\ \hat{F}_i^{-1}(0.9) \\ \$295,000 \\ \$1,246,000 \end{pmatrix}.$$

We then find the Pareto log-normal distribution which best matches these data. The Pareto log-normal distribution is characterized by three parameters, (α, ν, τ^2) . We find the parameters that minimize the sum of squared differences between the quantiles j_i and the values of the Pareto lognormal CDF at the incomes y_i . Letting j_{ik} denote the k^{th} element of j_i , and similarly for y_{ik} , the precise estimator we use is

$$(\alpha_i, \nu_i, \tau_i^2) = \arg \min_{(\alpha, \nu, \tau^2)} \sum_{k=1}^7 (j_{ik} - G^{-1}(y_{ik}; \alpha, \nu, \tau^2)).$$

B Proofs

Proof of Proposition 2. We begin by considering the case of pure career-switching and no intensive elasticity. First, note that the adjustment of marginal tax rates at any level must leave average externalities at each income level constant as the average entrant and exiter has the same externality as currently prevails at that income level. Second, note that the proposed tax satisfies the necessary conditions of Proposition 1, given that there is no intensive elasticity: for any y

$$\begin{aligned} E[\Delta T + \Delta E | \partial \Theta(y)] &= E[T(w_p h_p^*) + e_p w_p h_p^* - T(w_q h_q^*) - e_q w_q h_q^* | \partial \Theta(y)] = \\ E[E[T(y'') + e_p y'' - T(y') - e_q y' | \partial \Theta(y), w_q h_q^* = y', w_p h_p^* = y''] | \partial \Theta(y)] &= \\ E[E[T_0 - T_0 - E[e_{p^*} | \theta(y'')] y'' + E[e_{p^*} | \theta(y')] y' + e_p y'' - e_q y' | \partial \Theta(y), w_q h_q^* = y', w_p h_p^* = y''] | \partial \Theta(y)] &= 0. \end{aligned}$$

Next note that any other continuous tax scheme must violate these conditions. Suppose, to the contrary, that there is another scheme \hat{T} that obeys the conditions with $\hat{T} \neq T$. Then either $\hat{T} = T + k$ for some constant k or we can identify an open set $Y \subset (\underline{y}, \bar{y})$ such that $\hat{T}(y) - T(y) > \hat{T}(y') - T(y') \forall y \in Y$ and $y' \in (\underline{y}, \bar{y}) \setminus Y$. We deal with each of these cases separately:

1. $\hat{T} = T + k$: In this case, the allocation of every type to a profession is the same, but the tax scheme is either infeasible (if $k < 0$) or burns money ($k > 0$).
2. \hat{T} differs substantively: To show that the conditions of Proposition 1 are violated it is sufficient to show that any weighted sum of the conditions differs from 0. Because W is open it can be written as a countable union of n (where n possibly equals ∞) open intervals $(\underline{y}_1, \bar{y}_1), (\underline{y}_2, \bar{y}_2), \dots$. Consider

$$\begin{aligned} \sum_{i=1}^n E[\Delta T + \Delta E | \partial \Theta(\underline{y}_i)] f_S(\underline{y}_i) - E[\Delta T + \Delta E | \partial \Theta(\bar{y}_i)] f_S(\bar{y}_i) &= \\ \sum_{i=1}^n \int_{\partial \Theta(\underline{y}_i)} [\Delta T(\theta) + \Delta E(\theta)] f(\theta) d\theta - \int_{\partial \Theta(\bar{y}_i)} [\Delta T(\theta) + \Delta E(\theta)] f(\theta) d\theta &= \\ \sum_{i=1}^n \int_{\partial \Theta(\underline{y}_i) \setminus (\partial \Theta(\underline{y}_i) \cap \partial \Theta(\bar{y}_i))} [\Delta T(\theta) + \Delta E(\theta)] f(\theta) d\theta - \int_{\partial \Theta(\bar{y}_i) \setminus (\partial \Theta(\underline{y}_i) \cap \partial \Theta(\bar{y}_i))} [\Delta T(\theta) + \Delta E(\theta)] f(\theta) d\theta & \end{aligned}$$

Now note that

$$\partial\Theta(\underline{y}_i) \cap \partial\Theta(\bar{y}_i) = \left\{ \theta \in \Theta : \exists p, q \in 1, \dots, N : (w_p h_p^* < \underline{y}_i < \bar{y}_i < w_q h_q^*) \wedge (w_p h_p^* - T(w_p h_p^*) + \phi(h_p^*; \phi_p) = w_q h_q^* - T(w_q h_q^*) + \phi(h_q^*; \phi_q)) \right\},$$

which we abbreviate by $\partial\Theta_{y_p < \underline{y}_i < \bar{y}_i < y_q}$. On the other hand

$$\partial\Theta(\underline{y}_i) \setminus \left(\partial\Theta(\underline{y}_i) \cap \partial\Theta(\bar{y}_i) \right) = \partial\Theta_{y_p < \underline{y}_i < y_q < \bar{y}_i}$$

and

$$\partial\Theta(\bar{y}_i) \setminus \left(\partial\Theta(\underline{y}_i) \cap \partial\Theta(\bar{y}_i) \right) = \partial\Theta_{\underline{y}_i < y_p < \bar{y}_i < y_q}.$$

Thus all “switches” included are from incomes within Y to outside Y . Consider any income in $y \in (\underline{y}_i, \bar{y}_i)$ for some i (which is true for any $y \in Y$) and consider

$$E \left[\hat{T}(w_p h_p^*) + e_p w_p h_p^* - \hat{T}(y) - e_q y \mid w_q h_q^* \in Y, \theta \in \partial\Theta(\underline{y}_i) \cup \partial\Theta(\bar{y}_i) \right] <$$

$$T(w_p h_p^*) - T(w) + E \left[e_p w_p h_p^* - e_q y \mid w_q h_q^* \in Y, \theta \in \partial\Theta(\underline{y}_i) \cup \partial\Theta(\bar{y}_i) \right] = 0,$$

where the inequality follows by the definition of Y and the equality by the fact that, as above, changing the tax scheme does not change the externality properties of T as the average externality is static as individuals switch careers. Thus

$$\sum_{i=1}^n E \left[\Delta T + \Delta E | \partial\Theta(\underline{y}_i) \right] f_S(\underline{y}_i) - E \left[\Delta T + \Delta E | \partial\Theta(\bar{y}_i) \right] f_S(\bar{y}_i) < 0$$

contradicting the necessary conditions and establishing sufficiency.

Second we consider the case of pure intensive margin elasticities. In this case Proposition 1 immediately implies, given that the absence of correlation between e_{p^*} and $\epsilon_{p^*}^h$ implies the absence of covariance, that

$$\frac{([E[e_{p^*} | \Theta(y)] + T'(y)] E[\epsilon_{p^*}^h | \Theta(y)]) f(y)}{1 - T'(y)} = 0 \implies E[e_{p^*} | \Theta(y)] = -T'(y),$$

as $f(y), E[\epsilon_{p^*}^h | \Theta(y)] > 0$. This is also sufficient because hours decrease with taxes, establishing the necessary concavity. □

Proof of Proposition 4. By the same argument we used in the proof of Proposition 3, we need only to show that the sign of $\text{Cov}(e_{p^*(\theta)}, \epsilon_{p^*(\theta)}^h | \Theta(y))$ is the same as the sign of $e_1 - e_2$ at every income

level y . This is equivalent to showing that, for every income level y ,

$$E [\epsilon_1^h (\boldsymbol{\theta}) | \Theta(y), p^* (\boldsymbol{\theta}) = 1] > E [\epsilon_2^h (\boldsymbol{\theta}) | \Theta(y), p^* (\boldsymbol{\theta}) = 2] .$$

Recall that, from equation (3), that conditional on income being equal to y ,

$$\begin{aligned} \epsilon_i^h (\boldsymbol{\theta}) = \epsilon_i^h (w_i, \psi_i) &= \frac{w_i^2 [1 - T'(y)]}{h_i^* (w_i, \psi_i) [\phi'' (h_i^* (w_i, \psi_i)) + w_i^2 T''(y)]} = \frac{[1 - T'(y)] y^2}{(h_i^*)^3 \left[\phi'' (h_i^*) + \frac{y^2}{(h_i^*)^2} T''(y) \right]} = \\ &= \frac{[1 - T'(y)] y^2}{h_i^* [(h_i^*)^2 \phi'' (h_i^*) + y T''(y)]} , \end{aligned}$$

because, conditional on income y , $w_i = \frac{y}{h_i^*}$. Note that this expression clearly decreases in h_i^* . Thus we can rewrite the condition we are seeking to prove as

$$E [\epsilon_1^h (h_1^* (\boldsymbol{\theta})) | \Theta(y), p^* (\boldsymbol{\theta}) = 1] > E [\epsilon_2^h (h_2^* (\boldsymbol{\theta})) | \Theta(y), p^* (\boldsymbol{\theta}) = 2] .$$

If we can show that the y -conditional distribution of h_1^* first-order stochastically dominates the y -conditional distribution of h_2^* . Conditional on income y , we must have $h_i^* (w_i, \psi_i) w_i = y$. Note that both h_i^* and w_i are increasing in w_i and h_i^* is strictly increasing in w_i . Thus, conditional on income y , ψ_i is a decreasing function of w_i , while hours $h_i^* = \frac{y}{w_i}$ decrease in w_i as well. Thus hours must be an increasing function of ψ_i , conditional on income y . So we can rewrite our desired condition as

$$E [\epsilon_1^h (h_1^* (\psi)) | \Theta(y), p^* (\boldsymbol{\theta}) = 1] > E [\epsilon_2^h (h_2^* (\psi)) | \Theta(y), p^* (\boldsymbol{\theta}) = 2] ,$$

which holds because the income-contingent distribution of income in profession 1 dominates that in profession 2 and the expectation of any increasing function of a random variable which first-order stochastically dominates another must be greater than the expectation of the same function of the other random variable. \square