

Technical Change, Wage Inequality, and Taxes[†]

By LAURENCE ALES, MUSAB KURNAZ, AND CHRISTOPHER SLEET*

This paper considers the normative implications of technical change for tax policy design. A task-to-talent assignment model of the labor market is embedded into an optimal tax problem. Technical change modifies equilibrium wage growth across talents and the substitutability of talents across tasks. The overall optimal policy response is to reduce marginal income taxes on low to middle incomes, while raising those on middle to high incomes. The reform favors those in the middle of the income distribution, reducing their average taxes while lowering transfers to those at the bottom. (JEL D31, H21, H23, H24, J31, O33)

Technical change is inherently redistributive, complementing the labor of some whilst substituting for that of others. A large positive literature has analyzed its impact on the wage distribution. This literature has emphasized skill-biased technical change that favors the skilled over the unskilled and, more recently, has stressed the role of technical change in replacing “routine labor” in the middle of the wage distribution. However, while the positive literature documenting the redistributive nature of technical change is extensive, normative work exploring the policy implications of such change is not.¹ Our paper fills this gap. We explore how more than 30 years of technical change in the United States has affected the policy recommendations that economic theory provides. Overall, we find that such change creates a rationale for a modest adjustment of optimal policy in a direction that favors middle income earners, reducing their average taxes while lowering transfers to those at the bottom. Optimal marginal taxes are reduced on incomes that are low (but not the lowest) and raised on incomes that are high (but not the highest). Although, the

*Ales: Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213 (e-mail: ales@cmu.edu); Kurnaz: Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213 (e-mail: mkurnaz@andrew.cmu.edu); Sleet: Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213 (e-mail: csleet@andrew.cmu.edu). We thank V. V. Chari, Jonathan Heathcote, Sebastian Koehne, Hakkı Yazici, Marco Bassetto, Giovanni Gallipoli, Florian Scheuer, and many seminar participants for very useful comments and discussions. We thank Ben Tengelsen for excellent research assistance. Previous versions of this paper circulated under the title “Tasks, Talents and Taxes.” The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20140466> to visit the article page for additional materials and author disclosure statement(s).

¹For a historical account of the relationship between skill and technology see Goldin and Katz (1998), Autor, Katz, and Krueger (1998), and the references therein. Bresnahan, Brynjolfsson, and Hitt (2002) look at firm level evidence connecting technology and the demand for skills. Autor, Levy, and Murnane (2003) argue that recent technical change has led to the replacement of “routine” labor in the middle of the wage distribution. Autor, Katz, and Kearney (2006) and Goos and Manning (2007) document “job polarization”: growth in low- and high-skill occupations.

overall effects are moderate, they are the net effect of larger countervailing forces stemming from technical change. First, such change directly modifies wage differentials across differently talented workers; second it alters the substitutability of talent across occupations and, hence, the sensitivity of wage differentials to taxes. The evolution of optimal policy depends upon the balance of these conflicting forces.

We make theoretical and quantitative contributions. On the theoretical side, we embed a talent-to-task assignment model into an optimal tax framework. The former has been used by labor and trade economists to analyze the implications of technical change for the structure of wages and employment. We show how the technological parameters emphasized in this work shape optimal tax formulas. On the quantitative side, we bring a parametric assignment model to the data; we estimate the key parameters and derive the implications of technical change from the 1970s to the present day for policy.

The normative tax literature largely focuses on the incentive to supply effort by perfectly substitutable and privately informed workers. An exception is Stiglitz (1982) who allows for imperfect substitutability between the effort of two different talents.² This assumption renders relative wages sensitive to the profile of effort across talents and, hence, tax policy. In particular, Stiglitz identifies a wage compression motive for subsidizing high and taxing low talents. By doing so the wages of high talents are compressed relative to low and the former's incentive constraints are relaxed. We begin our analysis with a Stiglitz-type environment in which the production function is defined directly over the imperfectly substitutable labor input of many different worker types. In this setting with minimal restriction on the production function, we derive a general formula for optimal taxation. The formula provides a framework for interpreting subsequent results. Stiglitz's (1982) wage compression channel remains operative, but now takes a more complex form: the motive to tax a given talent type k at the margin depends, in part, on the elasticity of the relative wages of all pairs of adjacent talent types (ordered by wages) with respect to k 's effort. This setting suggests two ways in which technical change can influence optimal policy. First, factor augmenting technical change that is biased toward a subset of talents can do so by modifying relative wages and, hence, tightening or relaxing incentive constraints. Second, technical change that alters the effect of one talent type's effort on the relative wages of other talent types impacts policy by strengthening or diluting the wage compression channel described above.

We next embed an assignment model into an optimal tax framework.³ In the class of assignment models we consider, talented workers have a comparative advantage in complex tasks and assortative matching of workers to tasks occurs. To such

²Other important exceptions include Lockwood, Nathanson, and Weyl (2014); Rothschild and Scheuer (2013); Rothschild and Scheuer (2014); Rothschild and Chen (2014); and Slavík and Yazıcı (2014).

³The assignment framework originated with Roy (1950). Versions with a continuum of tasks, single dimensional talent, and comparative advantage of talented workers in complex tasks were developed by Sattinger (1975) and Teulings (1995). Such models have proven to be a rich laboratory for analyzing the role of task-talent distributions and the productivity of task-talent matches in shaping the wage distribution. Recently, these models have been used to explore the implications of technical change that attaches to tasks (rather than talents), see Costinot and Vogel (2010); Acemoğlu and Autor (2011); and Autor and Dorn (2013).

models we add an intensive effort margin, a societal motive for redistribution and explicitly private talent.⁴

In the equilibria of our embedded model, workers sort themselves efficiently across tasks *conditional on the effort of other workers*. This induces an indirect production function over the effort of different talents of the sort that our earlier analysis directly assumed. Technological parameters that determine relative task demand and the productivity of task-talent matches in the assignment framework are thus mapped to the variables and elasticities necessary for optimal tax analysis. In particular, the pattern of comparative advantage of talents across tasks shapes the sensitivity of relative wages to variations in the effort profile and, hence, policy. A local reduction in marginal taxes that induces a given talent type to increase its effort, depresses the (shadow) price of the task to which the type is assigned and, hence, the type's relative wage. Workers of this type offset this reduction by migrating into neighboring tasks, mitigating the impact on their original task's shadow price. However, the offset is partial since this migration erodes their productivity relative to neighboring talents. The greater is the comparative advantage of talented workers in complex tasks, the greater this erosion and the more sensitive are relative wages to task assignment. Thus, technical change that raises talent-complexity comparative advantage enhances the policymaker's ability to influence the wage structure through taxation. It strengthens the wage compression force identified in the more reduced form Stiglitz setting.⁵

We take our model to the data and quantify the implications of 30 years of technical change in the United States for optimal policy. We treat information on occupations, incomes, and hours in the Current Population Survey (CPS) as if it was generated by an equilibrium of our assignment model and use parametric assumptions and equilibrium restrictions to recover estimates of key technological parameters for the 1970s and the 2000s. To relate empirical occupations to the ordered set of tasks in our model, we order the former by the average wage paid. We recover an empirical proxy for the assignment of tasks to talents from the distribution of workers across occupations (ordered by wages). The estimation of parameters determining the demand for tasks is separated from those determining the productivity of task-talent matches by assuming a Cobb-Douglas technology for final goods as a function of tasks. This enables us to identify the demand parameters with occupational compensation shares. Parameters determining the productivity of talent-task matches and, hence, comparative advantage are recovered from the empirical assignment function and the distribution of wages across tasks using the envelope condition for wages implied by the model. After obtaining these estimates and supplementing them with calibrated preference parameters, we calculate optimal tax policies for the 1970s and 2000s.

⁴Rothschild and Scheuer (2013) were the first to consider the optimal tax implications of an assignment model. They do so in the context of a Roy model, i.e., a model with two sectors and no explicit notion of comparative advantage. We elaborate on the differences between their focus and approach and ours below.

⁵Migration of workers into neighboring tasks depresses the shadow prices of these tasks inducing the talents occupying them to migrate as well. A ripple effect is created and, so, an adjustment in one talent type's effort can induce reassignment of many types, affecting their relative wages and in the process relaxing and tightening many incentive constraints. However, the greater is talent-complexity comparative advantage the more contained the impact of a policy-induced effort adjustment.

We find evidence of relative reductions in demand for mid-level tasks and relative increases in demand for low- and high-level tasks. We also find evidence of a twisting of the talent-task productivity function, with low-talent productivity catching up to high talent in simple tasks and falling behind in more complex ones. The latter is associated with significant increases in the comparative advantage of more talented workers in more complex tasks. Moving from the 1970s to the 2000s, we find that under our benchmark estimation/parameterization, optimal marginal tax rates rise at the very bottom of the income distribution, fall on low- to middle-level incomes, rise on higher ones before falling again at the very top of the income distribution. This change in policy favors those in the middle of the income distribution who pay lower average taxes; optimal transfers to workers at the first and second income deciles are reduced. The twisting of the productivity function is the main force at work. It has two effects. First, it suppresses wage variation at the bottom of the income distribution, while enhancing it at the top. This relaxes incentive constraints on low incomes, while tightening them on high ones; it is a force for reductions in optimal marginal taxes on the former and increases on the latter. These effects are slightly enhanced by the relative reduction of demand for mid-level tasks populated by mid-level talents. Second, there is a partially offsetting strengthening of the wage compression channel. Higher comparative advantage of talented workers in complex tasks increases the policymaker's motive to apply high marginal taxes on low talents. Such taxes deter low-talent effort, raise low-level task prices, and encourage higher talents into these tasks. The relative productivity of these task migrants is eroded, suppressing their wage premia and relaxing incentive constraints. A parallel strengthening of the policymaker's motive to reduce marginal taxes on high talents occurs. Policy depends on the balance of these two forces. The first dominates at most incomes under our benchmark parametrization (except those in the extreme tails), but since the second dampens the first, the overall effect is modest.

The equilibrium of our baseline model does not exhibit intra-task wage dispersion or the payment of the same wage in multiple tasks ("wage overlap"). Thus, it cannot capture the policy implications stemming from these. At the end of the paper, in Section VI (with details and elaboration in online Appendix VI), we describe an extension that permits nondegenerate and overlapping supports for intra-task wage distributions. This extension incorporates a second talent dimension, which impacts absolute advantage alone. We find that our results concerning the implications of technical change for policy are qualitatively robust to, but quantitatively dampened by this extension. We use it to obtain a lower bound on the responsiveness of policy to technical change.

The remainder of the paper proceeds as follows. After a brief literature review, Section I provides motivating facts. Section II gives optimal tax formulas for economies with imperfectly substitutable labor types and provides an initial discussion of the implications of technical change for policy. In Section III, an assignment model is embedded into an optimal tax framework. An indirect production function over worker effort is derived and the parameters of the assignment model are related to the relevant terms of the optimal tax formulas from Section II. In addition, the implications of technical change for policy in a simple two talent model are discussed. Section IV describes how the assignment model is used to identify estimates of technical change and reports these estimates. In Section V, optimal policy for the

1970s and 2000s is computed and the implications of technical change for policy recovered. The tax formula from Section II is used to decompose and account for changes to optimal taxes. Section VI describes a model extension that can accommodate intra-task wage variation; Section VII concludes. Appendices contain proofs, robustness checks, and extensions.

Literature.—A contribution of our paper is to bring together the normative optimal taxation literature and a positive literature that analyzes the impact of technical change on the wage distribution. Both literatures are large. Many contributions to the latter have attributed increases in the skill premium to skill-biased technical change, formalizing this insight in what Acemoğlu and Autor (2011) have called the “canonical model,” i.e., a model with imperfectly substitutable skilled and unskilled workers and factor-augmenting technical change directed toward the skilled.⁶ Recently, a more nuanced view of the labor market has emerged that emphasizes growth in low- and high-wage occupations relative to those in the middle. It has spurred the development of assignment models that endogenize the joint distribution of workers across wages and occupations and in which technical change attaches to tasks rather than worker types. Examples include Acemoğlu and Autor (2011) and Autor and Dorn (2013).

Most contributions to the normative literature focus on the incentive to supply labor in environments with privately known talent and perfectly substitutable labor. Stiglitz (1982) was the first to introduce imperfectly substitutable labor into such a setting. Rothschild and Scheuer (2013) (extended in Rothschild and Scheuer 2014) were the first to introduce assignment. They show that a worker’s ability to select her task mutes the regressivity of optimal taxes found by Stiglitz. They also show that optimal tax formulas are substantially complicated by additional terms stemming from wage overlap. The focus in Rothschild and Scheuer (2013) is on economies with two tasks and two dimensional talents.⁷ In contrast, our baseline assignment model features a continuum of tasks and one dimensional talent. In our model a more talented worker is better at everything, but especially good at some things, with those things interpreted as more complex tasks. The restriction to one dimensional talents follows a tradition in labor economics initiated by Sattinger (1975) and adopted recently by the positive literature described above. Its adoption allows us to make contact with these recent contributions, to formulate the notions of talent and task complexity in a parsimonious way and to develop a strategy for bringing our model to the data. It permits a significant simplification of the tax formula in Rothschild and Scheuer (2013) (via the omission of wage overlap) and leads us to adopt a substantially different approach to analyzing the problem than that in Rothschild and Scheuer (2013). However, it cannot accommodate intra-task wage dispersion or wage overlap. In Section VI and online Appendix VI we provide an extension that can. Lockwood, Nathanson, and Weyl (2014) also integrate tax considerations into an assignment setting. They focus on the externalities associated with certain

⁶Examples include Acemoğlu (2002) and Krusell et al. (2000).

⁷This is generalized in Rothschild and Scheuer (2014) to K tasks and K dimensional talent.

assignments and characterize the optimal structure of corrective Pigouvian taxation.⁸ We abstract completely from this tax motive.

Slavík and Yazici (2014) apply the logic of Stiglitz (1982) to capital taxation. In their paper they introduce two sorts of capital: buildings and machines. Following the skill premium literature, they assume a machine-skill (or machine-talent) complementarity. Thus, machines raise the marginal product of the talented relative to the untalented and, as in Stiglitz (1982), this dilutes incentives. It is socially desirable to deter the accumulation of machines. In quantitative work, Slavík and Yazici (2014) show that this creates a rationale for quite high rates of (machine) capital taxation. Slavík and Yazici's (2014) contribution is complementary to ours. They endogenize technical change in the context of a two talent "canonical model" and develop policy implications. We treat technical change parametrically, but do so in a multi-talent/multi-task assignment setting.

Heathcote, Storesletten, and Violante (2014) analyze optimal income tax progressivity in a rich dynamic environment. They assume imperfectly substitutable skills, but do not explicitly model tasks. Our model is static, but we add assignment and, hence, endogenize the substitutability of skills and relate it to technical change. In addition, Heathcote, Storesletten, and Violante (2014) restrict optimal taxes to a parametric class, we do not.

I. Evolution of the Occupational Wage Distribution: Stylized Facts

We first document some stylized facts that motivate our analysis. Figure 1 displays changes in average incomes across (one-digit) occupations from the 1970s to the present.⁹ The figure indicates considerable variation in the experience of different occupations, with some exhibiting significant average income growth and others stagnating. Moreover, occupations with slow average income growth were predominantly middle income in the 1970s, while fast growers were mainly low or high income at that time. For example, precision production, craft, and repair workers had a mid-level income of \$33,109 in 1975 (all incomes are expressed in 2005 dollars) and negligible income growth subsequently. In contrast, the two occupations with the fastest growing average incomes, services and managerial and professional, had average incomes in the mid-1970s of \$12,912 and \$40,013, placing them at opposite extremes of the income distribution. Such occupational polarization, with the middle growing more slowly than the extremes, is not confined to earnings; it is also present in various measures of occupational size and demand. Figure 2 displays changes in the share of employment of different occupations over time.¹⁰

Here managerial/professional and service related occupations that are concentrated in the extremes of the income distribution are expanding in size, while mid-income level occupations such as operators and fabricators (mostly employed in manufacturing) whose workers are shrinking over time.

⁸Rothschild and Scheuer (2014) also incorporate this motive into their theoretical work.

⁹The data is taken from the March CPS. See online Appendix III for additional details on the data and our sample selection.

¹⁰See, *inter alia*, Goos and Manning (2007); Acemoglu and Autor (2011); and Autor and Dorn (2013) for related evidence.




FIGURE 1. EVOLUTION OF AVERAGE INCOME BY OCCUPATION OVER TIME

Source: March CPS.




FIGURE 2. EVOLUTION IN SIZE OF EMPLOYMENT BY OCCUPATION OVER TIME

Source: March CPS.

Overall, the picture that emerges from the CPS (and other data sources) is one in which high-wage and low-wage occupations are growing in size and in average compensation relative to middle ones. If talent is imperfectly substitutable across occupations, then these varied occupational fortunes suggest varied fortunes for differently-talented workers. In the remainder of the paper we consider the optimal policy response to such events.

II. Taxation with Imperfectly Substitutable Workers

Mirrlees's (1971) model of optimal taxation assumes that workers of different types are perfect substitutes and that final output is a weighted sum (or integral) of worker efforts, with the weights given by private productivities. Stiglitz (1982) allows for a more general production function. He assumes that workers are one of two imperfectly substitutable types and interprets these types as "low" and "high" skilled. In this section, we generalize (Stiglitz 1982) to K -types, but place no interpretation on a worker's type (the nature of which is defined implicitly by the production function). In this (and compared to later sections reduced form) context we discuss implications of technical change for taxes.¹¹

A. Physical Environment

Workers.—A continuum of workers has identical preferences over consumption $c \in \mathbb{R}_+$ and effort $e \in [0, \bar{e}]$ described by a utility function $U: \mathbb{R}_+ \times [0, \bar{e}] \rightarrow \mathbb{R}$. The function U is assumed to be concave, twice continuously differentiable on the interior of its domain, with for each $e \in [0, \bar{e}]$, $U(\cdot, e)$ increasing and for each $c \in \mathbb{R}_+$, $U(c, \cdot)$ decreasing and strictly concave. First and second partial derivatives of U are denoted U_x and U_{xy} with $x, y \in \{c, e\}$. U satisfies the Inada conditions: for all $c > 0$, $\lim_{e \downarrow 0} U_e(c, \cdot) = 0$ and $\lim_{e \uparrow \bar{e}} U_e(c, \cdot) = -\infty$. In addition, U satisfies the Spence-Mirrlees single crossing property: $-U_e(c, y/w) / \{wU_c(c, y/w)\}$ is decreasing in w . Workers are partitioned across a finite number of "types" $K \geq 2$ with a fraction π_k of workers being of type $k \in \{1, \dots, K\}$. The fraction of workers with type less than or equal to k is denoted $\Pi_k = \sum_{j=1}^k \pi_j$.

Workers sell their labor to firms and pay taxes on the income that they earn. Let $T: \mathbb{R}_+ \rightarrow \mathbb{R}$ denote an income tax function.¹² A worker of type k receiving wage w_k solves the problem

$$(1) \quad \max_{\mathbb{R}_+ \times [0, \bar{e}]} U(c, e) \quad \text{s.t.} \quad c \leq w_k e - T(w_k e).$$

Technology.—A representative competitive firm hires workers of all types. The firm uses a production function $F: \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ defined directly on the labor inputs of the different types. The firm solves

$$\max_{\mathbb{R}_+^K} F(e_1 \pi_1, \dots, e_K \pi_K) - \sum_{k=1}^K w_k \pi_k e_k,$$

¹¹ Much of the optimal tax literature is cast in terms of a continuum of types. This literature maintains the linear production function assumption. Although versions of the results that we give below are available for continuum economies, for *general* constant returns to scale production functions over worker effort allocations, their derivation requires leaving the framework of optimal control and maximizing an infinite-dimensional Lagrangian directly. To avoid technical complications that do not generate additional economic insight we do not do this. In our later assignment setting, optimal control techniques are applicable and we use them in online Appendix II.

¹² We restrict attention to non-stochastic tax functions. See Hellwig (2007) for sufficient conditions for such mechanisms to be socially optimal in utilitarian settings.

where e_k is the common effort level of workers of type k . F is assumed to be a continuously differentiable, constant returns to scale function with k th partial derivative F_k . At this stage, we place no further restrictions on F . In classical Mirrlees models $F(e_1 \pi_1, \dots, e_K \pi_K) = \sum_{k=1}^K a_k e_k \pi_k$ for some positive constants $\{a_k\}$ and workers of different types are perfectly substitutable. However, we allow for and focus upon worker types that are imperfect substitutes in production. Since F defines what it means for a worker to be of one type or another, the economic nature of a worker's type is for the moment left implicit.¹³

Tax Equilibrium.—Let $G \in \mathbb{R}_+$ be a fixed public spending amount. Given G , a tax equilibrium is an income tax function $T: \mathbb{R}_+ \rightarrow \mathbb{R}$, an allocation $\{c_k, e_k\}_{k=1}^K$ and a wage profile $\{w_k\}_{k=1}^K$ such that (i) for each $k = 1, \dots, K$, (c_k, e_k) solves (1), (ii) for each $k = 1, \dots, K$, $w_k = F_k(e_1 \pi_1, \dots, e_K \pi_K)$, and (iii) the goods market clearing condition holds: $G + \sum_{k=1}^K c_k \pi_k \leq F(e_1 \pi_1, \dots, e_K \pi_K)$. Let \mathcal{E} denote the set of tax equilibria (given G), which we take to be nonempty.

B. Optimal Policy

A government attaches Pareto weight g_k to workers of type k , with weights normalized to satisfy $\sum_{k=1}^K g_k = 1$. It selects a tax equilibrium to solve

$$(PP) \quad \sup_{\mathcal{E}} \sum_{k=1}^K U(c_k, e_k) g_k.$$

Let T^* and $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$ denote an optimal tax equilibrium. Define the corresponding (optimal) marginal tax rate at income $q_k^* := w_k^* e_k^* > 0$ to be¹⁴

$$\tau_k^* = 1 + \frac{U_e(c_k^*, e_k^*)}{w_k^* U_c(c_k^*, e_k^*)}.$$

To characterize optimal tax equilibria, we follow the conventional procedure of recovering optimal allocations from a mechanism design problem. Subsequently, prices and (optimal) taxes are determined to ensure implementation of this allocation as part of a tax equilibrium. The mechanism design problem associated with (PP) is

$$(MDP) \quad \sup_{\{c_k, e_k\}_{k=1}^K \in \{\mathbb{R}_+ \times [0, \bar{e}]\}^K} \sum_{k=1}^K U(c_k, e_k) g_k$$

¹³The firm's problem determines relative levels of efforts across types. The scale of the representative firm is determined in equilibrium.

¹⁴(PP) does not uniquely determine T^* . However, T^* may be chosen to be directionally differentiable in which case: $\partial T^-(q_k^*) \leq \tau_k^* \leq \partial T^+(q_k^*)$, where $\partial T^-(q_k^*)$ and $\partial T^+(q_k^*)$ are left and right derivatives of T^* at $q_k^* > 0$. If T^* is (chosen to be) differentiable at q_k^* , then its derivative at that point equals τ_k^* .

s.t. for each $k, j \in \mathcal{K} := \{(l, m) \in \{1, \dots, K\}^2, l \neq m\}$,

$$(2) \quad \eta_{k,j}: \quad U(c_k, e_k) \geq U\left(c_j, \frac{F_j(e_1 \pi_1, \dots, e_K \pi_K)}{F_k(e_1 \pi_1, \dots, e_K \pi_K)} e_j\right)$$

and

$$(3) \quad \chi: \quad F(e_1 \pi_1, \dots, e_K \pi_K) \geq G + \sum_{k=1}^K c_k \pi_k.$$

In (MDP) the government selects a report-contingent allocation of consumption and effort $\{c_k, e_k\}_{k=1}^K$ that induces each worker to truthfully report its type k and produce the associated income $q_k = F_k(e_1 \pi_1, \dots, e_K \pi_K) e_k$. Incentive constraints that ensure the optimality of truthful reporting are given in (2) with corresponding Lagrange multipliers $\eta_{k,j}$. If type k claims to be of type j she must reproduce the corresponding income $q_j = F_j(e_1 \pi_1, \dots, e_K \pi_K) e_j$. The effort cost to her of doing so is $\frac{F_j(e_1 \pi_1, \dots, e_K \pi_K)}{F_k(e_1 \pi_1, \dots, e_K \pi_K)} e_j$. Thus, the (k, j) th incentive constraint (2) depends upon the *entire profile* of worker efforts via the (k, j) th shadow wage ratio. We refer to a (k, j) -incentive constraint as local if $j = k - 1$ or $j = k + 1$, local downwards if $j = k - 1$ and local upwards if $j = k + 1$. The final restriction (3) in (MDP) is the resource constraint with corresponding multiplier χ .

Toward understanding how technical change shapes policy, we give a proposition that relates optimal taxes to F . This proposition is a consequence of a more general result given in the Appendix. In the latter, we show that when worker types are ordered consistently with optimal wages and incomes, then only local downwards $(k, k - 1)$ or upwards $(k, k + 1)$ incentive constraints bind. In the main text we follow the common convention of assuming that only the former are binding and then verifying this assumption in numerical calculations as needed.¹⁵ To state the proposition (and its generalization in the Appendix) it is convenient to reexpress the constraints in (MDP) in the form $\mathcal{G}(\{c_k, e_k\}_{k=1}^K) \geq 0$, where $\mathcal{G}: \mathbb{R}_+^{2K} \rightarrow \mathbb{R}^{K(K-1)+1}$ combines the constraint functions. Problem (MDP) satisfies a (Mangasarian-Fromowitz) constraint qualification at $\{c_k, e_k\}_{k=1}^K \in \mathbb{R}_{++}^{2K}$ if there is an $x \in \mathbb{R}^{2K}$ such that $\nabla \mathcal{G}(\{c_k, e_k\}_{k=1}^K) x < 0$, where $\nabla \mathcal{G}(\{c_k, e_k\}_{k=1}^K)$ is the Jacobian of \mathcal{G} at $\{c_k, e_k\}_{k=1}^K$. Let $\eta_{k,j}^*$ and χ^* denote the optimal (Karush-Kuhn-Tucker) multipliers associated with the incentive and resource constraints. Finally, let $\Delta_e U_c(c', e'; \delta) := \frac{U_c(c', e' + \delta) - U_c(c', e')}{\delta}$ denote a finite difference approximation to the derivative of U_c with respect to e at (c', e') and define $\Delta_e U_e$ analogously.

PROPOSITION 1: *Let T^* and $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$ denote an optimal tax equilibrium with worker types indexed so that $w_k^* = F_k(e_1^* \pi_1, \dots, e_K^* \pi_K)$ is nondecreasing in k . Assume that for each k , $c_k^* \in \mathbb{R}_{++}$ and $e_k^* \in (0, \bar{e})$, that \mathcal{G} satisfies the constraint*

¹⁵ A general formula with possibly binding upwards incentive constraints is supplied in the Appendix.

qualification at $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$ and that the local upwards incentive constraints are nonbinding, i.e.,

$$(NUIC) \quad U(c_k^*, e_k^*) > U(c_{k+1}^*, q_{k+1}^*/w_k^*), \quad \text{where } q_{k+1}^* := w_{k+1}^* e_{k+1}^*.$$

Optimal tax rates then satisfy

$$(4) \quad \frac{\tau_k^*}{1 - \tau_k^*} = \underbrace{\frac{1 - \Pi_k}{\pi_k} \frac{\Delta w_{k+1}^*}{w_{k+1}^*} \Psi_k^* \mathcal{H}_k^*}_{\text{Mirrlees}} + \underbrace{\sum_{j=1}^{K-1} \mathcal{M}_{k,j}^* \phi_{k,j}^*}_{\text{Wage Compression}},$$

where $\Delta w_k^* := w_k^* - w_{k-1}^*$,

$$\Psi_k^* := \sum_{j=k}^{K-1} \mathcal{N}_{k,j}^* \left\{ 1 - \frac{g_{j+1} U_c(c_{j+1}^*, e_{j+1}^*)}{\pi_{j+1} \chi^*} \right\} \left(\frac{U_c(c_k^*, e_k^*)}{U_c(c_{j+1}^*, e_{j+1}^*)} \right) \frac{\pi_{j+1}}{1 - \Pi_k},$$

with $\mathcal{N}_{k,j}^* := \prod_{i=k+1}^j \frac{U_c(c_i^*, q_i^*/w_{i+1}^*)}{U_c(c_i^*, e_i^*)},$

$$\begin{aligned} \mathcal{H}_k^* := & - \frac{\Delta_e U_c(c_k^*, e_k^*; q_k^*/w_{k+1}^* - e_k^*)}{U_c(c_k^*, e_k^*)} e_k^* \\ & + \frac{\Delta_e U_e(c_k^*, e_k^*; q_k^*/w_{k+1}^* - e_k^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_j^*} e_k^* + 1, \end{aligned}$$

$$\mathcal{M}_{k,j}^* := \frac{U_c(c_k^*, e_k^*)}{U_e(c_k^*, e_k^*) e_k^*} \frac{U_e(c_j^*, q_j^*/w_{j+1}^*)}{U_c(c_j^*, e_j^*)} \frac{q_j^*}{w_{j+1}^*} \frac{1 - \Pi_j}{\pi_j} \frac{\pi_j}{\pi_k} \Psi_j^*$$

and cross relative wage elasticities

$$\phi_{k,j}^* := \frac{e_k^*}{w_{j+1}^*/w_j^*} \frac{\partial w_{j+1}^*/w_j^*}{\partial e_k} (e_1^*, \dots, e_K^*).$$

PROOF:

See the Appendix.

The right-hand-side of the optimal tax formula (4) is the sum of two terms, which we label “Mirrlees” and “Wage Compression.” We discuss each in turn below.

Mirrlees Term.—The Mirrlees term in (4) is quite standard in optimal tax analyses. We very briefly review and interpret its four components.¹⁶ \mathcal{H}_k^* is a discrete

¹⁶For detailed discussion of these components in a continuous-type setting see Salanié (2011).

approximation to $\frac{1 + \mathcal{E}_{u,k}}{\mathcal{E}_{c,k}}$, where $\mathcal{E}_{c,k}$ and $\mathcal{E}_{u,k}$ are, respectively, the compensated and uncompensated labor supply elasticities at (c_k^*, e_k^*) . If worker preferences are additively separable, this reduces to one plus (a discrete approximation to) the reciprocal of the Frisch elasticity. Incentive-compatibility considerations require that if worker type k receives an increment in consumption all higher types $j = k + 1, k + 2, \dots, K$ receive an increment in utility sufficient to deter them from reporting a lower type. Ψ_k^* captures the net societal cost of such a redistribution; it weighs the cost of extracting resources from the population at large against the benefits of raising the welfare of higher income types. $\frac{1 - \Pi_k}{\pi_k}$ is the reciprocal of the type hazard. This plays an important role in conventional optimal tax analysis since, if types have compact support (as in the current finite setting), it implies zero marginal taxes at the maximal income. However, it is unaffected by technical change and, thus, is less central to our analysis. In contrast, the wage growth (across types) term $\Delta w_{k+1}^*/w_{k+1}^*$ is endogenous and important in what follows. To understand its role consider the local downwards incentive constraint:

$$(5) \quad U(c_{k+1}, e_{k+1}) \geq U(c_k, e_k w_k / w_{k+1}).$$

As noted previously, the wage ratio w_{k+1}/w_k appears on the right-hand-side of this inequality. Higher values of this ratio reduce the effort that a $k + 1$ th type worker must exert to mimic a k th-type. Consequently, they tighten the incentive constraint and lead to greater distortions of allocations. Higher wage growth across the k and $k + 1$ types is, other things equal, a force for higher marginal taxes on the k th type.¹⁷

Wage Compression Term.—The second term in (4) does *not* appear in standard optimal tax equations that are derived from models with linear production functions and exogenous wages. In settings with nonlinear production functions, such as ours, the effort of the k th worker type can affect the marginal rate of transformation and, hence, the ratio of wages between the j and $j + 1$ th types. Following the logic of the previous paragraph, more compressed wage ratios relax incentive constraints and, thus, to the extent that the effort of a given type increases wage differentials it should be deterred. In particular, larger values of the cross relative wage-effort elasticities $\phi_{k,j}^*$ are a rationale for higher marginal taxes on the k th type's income. Stiglitz (1982) identifies this wage compression channel in a two type model. In that case there is only one binding incentive constraint and $-\phi_{1,1}^* = \phi_{2,1}^* = 1/\mathcal{E}^*$, where \mathcal{E}^* is the elasticity of substitution between the two worker types i.e.,

¹⁷The terms $\frac{1 - \Pi_k}{\pi_k}$ and $\frac{\Delta w_{k+1}^*}{w_{k+1}^*}$ may be consolidated as: $\frac{1 - \Pi_k}{\pi_k} \frac{\Delta w_{k+1}^*}{w_{k+1}^*}$. In the continuous limit the latter reduces to $\text{Haz}(w) = \frac{1 - \Xi(w)}{\xi(w)w}$, where Ξ and ξ are the wage distribution and density functions and, following the usage of Saez (2001), $\text{Haz}(w)$ is the wage hazard ratio. In the continuous setting, the impact of a change in wage growth across types $\frac{1}{w^*(k)} \frac{\partial w^*}{\partial k}(k)$ on marginal taxes may be understood via its impact on Haz . Specifically, an increase in $\frac{\partial w^*}{\partial k}(k)$ reduces $\xi(w^*(k))$ (the “fraction” who will be distorted by a marginal tax) relative to $1 - \Xi(w^*(k))$ (the fraction who will be undistorted and will pay higher average taxes). It is, therefore, a force for higher marginal taxes at $w^*(k)$.

$\left(\frac{w_2/w_1}{e_2/e_1} \frac{\partial e_2/e_1}{\partial w_2/w_1}\right)$ at the optimum. Assuming this is positive, compression of wages between the two types, requires that the effort of the high (resp. low) type should be relatively encouraged (resp. discouraged). Since the first term in (4) is zero for $k = K = 2$, this translates into an optimal marginal income subsidy for high types and an enhanced marginal income tax for low types.

The Role of F .—The functional form for F plays an important role in shaping wage growth across types $\Delta w_{k+1}^*/w_{k+1}^*$, the cross relative wage-effort elasticities $\phi_{k,j}^*$ and, hence, optimal taxes. Consider first some standard functional forms. If F is a weighted sum of type efforts, as in the classical Mirrlees model, then $\Delta w_{k+1}^*/w_{k+1}^*$ is treated as structural and invariant to policy, while each $\phi_{k,j}^*$ is set to zero and the wage compression channel is shut down. If F is a more general CES function, then, except in the limiting linear case, policy can affect $\Delta w_{k+1}^*/w_{k+1}^*$ and the wage compression channel is active. However, the elasticity of substitution and, hence, the relative wage-effort elasticities remain parametric and independent of policy. The CES form also places strong restrictions on the latter requiring that they equal

$$(6) \quad \phi_{k,j}^* = \begin{cases} -\frac{1}{\mathcal{E}} & j = k \\ \frac{1}{\mathcal{E}} & j = k - 1 \\ 0 & \text{otherwise,} \end{cases}$$

where \mathcal{E} is the elasticity of substitution between the effort of worker type pairs. Thus, for each worker type k , the elasticities $\phi_{k,j}^*$ are nonzero only locally (i.e., a variation in a type's effort only affects its wage relative to others, it does not affect the relative wage of other type pairs) and all elasticities $\phi_{k,k}^*$ and $\phi_{k,k-1}^*$ take common values independent of k . These features have led to some resistance amongst labor and public finance economists to the use of CES production functions in modeling labor demand. For example, Salanié (2011, p. 111) asserts: "It is, unfortunately, quite difficult to specify a production function that models the limits to factor substitution with an infinite number of factors." He emphasizes that the substitutability of similar and dissimilar worker types may be quite different, but that such differences cannot be accommodated under the CES assumption. In Section III, we use an assignment framework to micro-found a production function F defined directly over worker efforts. In this setting the allocation of workers to tasks is efficient given worker effort and F is the upper envelope to a family of production functions indexed by worker task choices. Relative wage-effort elasticities $\phi_{k,j}^*$ and the wage compression channel are then shaped by both technology and policy and tied to the underlying technological parameters that describe the productivity of differently talented workers in different tasks.¹⁸

¹⁸Rothschild and Scheuer (2013) and its generalization Rothschild and Scheuer (2014) consider assignment economies in which workers choose tasks as well as effort. In their setting a worker's multidimensional type gives his or her productivity in all tasks. If, in the latter paper, the distribution over worker types places all mass on types that have positive productivity in only one task, then it reduces to a K -type Stiglitz economy with workers effectively "locked" into particular tasks. Consequently, results similar to Proposition 1 would emerge in Rothschild and

Technical Change.—The formulas in Proposition 1 point to several channels through which technical change can influence optimal policy. If such change raises the return to effort of all workers equally at a given effort profile, then it does not directly affect wage growth over types $\Delta w_{k+1}^*/w_k^*$ or the responsiveness of relative wages to effort $\phi_{k,j}^*$ in (4). Such “type-neutral” technical change impacts policy only insofar as it affects labor supply elasticities and relative marginal rates of substitution across workers.¹⁹ If, on the other hand, technical change augments the effort of a subset of workers, then, in general, it does affect wage growth over types. Specifically, if F is a CES function of the form $F(e_1 \pi_1, \dots, e_K \pi_K)$

$$= A \left[\sum_{k=1}^K D_k e_k^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \text{ then:}$$

$$(7) \quad \frac{\Delta w_{k+1}}{w_{k+1}} \approx -\log\left(\frac{w_k}{w_{k+1}}\right) = -\log\left(\frac{D_k}{D_{k+1}}\right) + \frac{1}{\varepsilon} \log\left(\frac{e_k}{e_{k+1}}\right),$$

and technically induced variations in the log relative CES weights $\left\{ \log \frac{D_k}{D_{k+1}} \right\}$ additively translate the map from efforts to wage growth over types. Such variations, by modifying the productivity of one type of worker relative to another at a given effort profile, relax or tighten incentive constraints and, hence, elicit an optimal tax response. They do not affect the responsiveness of relative wages to effort, i.e., they leave the elasticities $\phi_{k,j}$ unaltered (at the fixed values given in (6)). For more general production functions (such as the induced F in the next section), technical change can influence the sensitivity of wages to the effort profile as well. In particular, by reducing substitutability between skills, technical change can enhance the impact of variations in relative labor supplies on relative wages and, hence, the policymaker’s influence over the wage distribution. This strengthens the wage compression motive and is a further channel via which technical change can influence optimal policy (and a channel that is absent under the CES specification).

III. Taxation, Assignment, and Technical Change

We now consider optimal taxation in a framework with task assignment. As noted in the introduction, assignment-based frameworks have been used in the positive literature to formalize the impact of technical change on the distribution of workers across wages and occupations. As we show below they imply and, hence, micro-found an indirect production function over worker efforts. Consequently, we are able to relate key elasticities in the optimal tax equation (4) to deeper structural parameters that describe the relative demand for tasks and the way in which tasks and talent interact. We interpret changes in these parameters as technical change and conclude this section by deriving implications of such change for optimal policy in a very simple assignment model.

Scheuer (2014) under this restriction. Our analysis in Section III, however, shows that Proposition 1 is much more generally applicable.

¹⁹In our later numerical work, we shut this channel down by restricting to utility functions: $\log c + h(e)$.

A. Physical Environment

As before, workers are partitioned across types $1, \dots, K$ with a fraction π_k being of type k . Types are now explicitly identified with talents. In addition, there is a continuum of tasks $v \in [\underline{v}, \bar{v}]$ differentiated by complexity. Workers can choose which task to work (exert effort) in; they cannot work in multiple tasks. They face a schedule of task-specific wages $\omega : [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$, with $\omega(v)$ the wage per unit of effective labor paid in task v . A worker of talent k has productivity $a_k(v) \in \mathbb{R}_+$ in task v . If she exerts effort e in this task her effective labor is $a_k(v)e$ and her income is $\omega(v)a_k(v)e$. The worker chooses her consumption, effort, and task to solve

$$(8) \quad \sup_{\mathbb{R}_+ \times [0, \bar{e}] \times [\underline{v}, \bar{v}]} U(c, e) \quad \text{s.t.} \quad c \leq \omega(v)a_k(v)e - T(\omega)a_k(v)e.$$

The productivity functions $\{a_k\}$, $a_k : [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$, play a key role in the subsequent analysis. The following condition is imposed upon them.

ASSUMPTION 1: *The functions $a_k : [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$, $k \in \{1, \dots, K\}$ are continuous and satisfy (i) (weak comparative advantage) for each $k \in \{1, \dots, K-1\}$ and $v', v \in [\underline{v}, \bar{v}]$ with $v' > v$, $\log a_{k+1}(v') - \log a_k(v') \geq \log a_{k+1}(v) - \log a_k(v)$ and (ii) (absolute advantage) for each $k \in \{1, \dots, K-1\}$, $a_{k+1} > a_k$.*

By Assumption 1(i) a is a weakly log super-modular function of talent and task and higher talents have a weak comparative advantage in more complex tasks. In the subsequent analysis this assumption is often strengthened to strict log super-modularity: for $k \in \{1, \dots, K-1\}$, and $v', v \in [\underline{v}, \bar{v}]$ with $v' > v$, $\log a_{k+1}(v') - \log a_k(v') > \log a_{k+1}(v) - \log a_k(v)$. This stronger condition ensures assortative matching of tasks and talents in equilibrium. Assumption 1(ii) implies that more talented types have an absolute advantage in all activities. It is not essential for all of our results, but it guarantees that wages are strictly increasing in talent. Hence, the orderings over talent and wages conform and there is no “wage pooling” (multiple talents earning the same wage).

REMARK 1 (Interpreting a): *The function a captures the idea that different workers may be more or less effective at performing specific tasks or using task-specific capital. Combined with Assumption 1 it formalizes the notions of talent and task complexity. More talented workers are better at all tasks and are especially good at more complex ones. Relatedly, more complex tasks are more talent-intensive. The formulation of production here follows that in the assignment literature, e.g., Costinot and Vogel (2010), with the important addition of an intensive effort margin.²⁰*

²⁰The assignment literature refers to a worker's innate productive attribute as “skill.” Since skills are endogenous, we prefer the word talent. Our model could be reinterpreted as one in which workers exert effort partly or wholly in acquiring skills rather than working.

Later we allow for the possibility that a may change over time. We interpret such change as technical progress and allow it to depend upon both worker talent and task complexity. In particular if, for each v and $k' > k$, $\log \frac{a_{k'}(v)}{a_k(v)}$ increases, then technical progress is talent-biased; if for each k and $v' > v$, $\log \frac{a_k(v')}{a_k(v)}$ increases, then it is complexity-biased and if for each $k' > k$, $v' > v$, $\log \left(\frac{a_{k'}(v')}{a_k(v')} / \frac{a_k(v)}{a_k(v)} \right)$ increases, then it is biased toward high-talent, high-complexity matches. In the latter case, it enhances the comparative advantage of talent in complex tasks and reduces the substitutability of talent across tasks.

The task choices of workers imply a distribution of workers and, hence, effective labor across tasks. Let Λ_k denote a distribution of k th talent workers over tasks with density λ_k . If k th talent workers exert effort e_k , then the supply of effective labor in task v is

$$\sum_{k=1}^K \lambda_k(v) a_k(v) e_k.$$

A representative firm hires effective labor to perform tasks and combines task output to produce final output. Let $l: [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$ denote an allocation of effective labor across tasks and let \mathcal{L} denote the set of such allocations (with \mathcal{L} restricted to ensure the integrals defined below in (9) are well-defined). Output is assumed to equal effective labor in each task v . Final output Y is produced from task output and, hence, from an allocation of effective labor l using a CES-technology:

$$(9) \quad Y = H(l) := \begin{cases} A \left\{ \int_{\underline{v}}^{\bar{v}} b(v) l(v)^{\frac{\varepsilon-1}{\varepsilon}} dv \right\}^{\frac{\varepsilon}{\varepsilon-1}} & \varepsilon \in \mathbb{R}_+ \setminus \{1\}, \\ A \exp \left\{ \int_{\underline{v}}^{\bar{v}} b(v) \ln l(v) dv \right\} & \varepsilon = 1, \end{cases}$$

where $A > 0$ and $b: [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_{++}$ is a continuous function such that if $B(v) := \int_{\underline{v}}^v b(v') dv'$, then $B(\bar{v}) = 1$. Let $\omega: [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$ be the wage per unit of effective labor in each task v . The firm solves

$$(10) \quad \max_{l \in \mathcal{L}} H(l) - \int_{\underline{v}}^{\bar{v}} \omega(v) l(v) dv.$$

REMARK 2 (Interpreting b): The function b weights task output in the final good aggregator. Variations in b may be interpreted as stemming from technological or preference-based variations in demand for different task outputs. We do not explicitly model capital. However, the model may be extended in this direction, in which case the production functions in (9), under the assumption $B(\bar{v}) \in (0, 1)$, can be reinterpreted as indirect production functions for labor across tasks after the substitution of optimal capital. The parameter $b(v)$ is then interpreted as the sensitivity of final output with respect to the labor input in task v . It is influenced not only by variations in demand for different tasks, but also variations in the capital/labor intensity of tasks. Such variations are stressed by Acemoglu and Autor (2011) who emphasize

the automatization of middle complexity tasks. A further possibility is that b captures the extent to which workers purchase task output in domestic markets, produce it at home or purchase it in foreign markets. Shifts in b for some tasks may reflect the substitution of market for home production as in Buera and Kaboski (2012) or domestic for foreign production as in Grossman and Rossi-Hansberg (2008).

B. Tax Equilibria and the Government's Policy Problem

In the assignment setting, the definition of a tax equilibrium is modified as follows.²¹

Tax Equilibrium.—Let G be a fixed public spending amount. Given G , a *tax equilibrium* is an income tax function $T: \mathbb{R}_+ \rightarrow \mathbb{R}$, an allocation $\{l, \{c_k, e_k, \lambda_k\}_{k=1}^K\}$ and a wage profile ω such that (i) for each $k = 1, \dots, K$, (c_k, e_k) and v in the support of Λ_k solves the k th worker's problem (8) at T and ω ; (ii) l solves (10) at ω ; (iii) the final goods market clears:

$$(11) \quad G + \sum_{k=1}^K c_k \pi_k \leq H(l);$$

and (iv) labor markets clear, for all $v \in [\underline{v}, \bar{v}]$,

$$(12) \quad l(v) = \sum_{k=1}^K \lambda_k(v) a_k(v) e_k,$$

and for all $k = 1, \dots, K$,

$$(13) \quad \pi_k = \int_{\underline{v}}^{\bar{v}} \lambda_k(v) dv.$$

Again, let \mathcal{E} denote the set of tax equilibria. Proposition 2 below characterizes tax equilibria. It contains the simple, but important result that *conditional on effort* assignment in a tax equilibria maximizes output.

PROPOSITION 2: *Let Assumption 1 hold. Let $\{l, \{c_k, e_k, \lambda_k\}_{k=1}^K\}$ and ω be, respectively, the allocation and wage profile of a tax equilibrium. Then there is a tuple of threshold tasks $\{\tilde{v}_k\}_{k=1}^{K-1}$ such that:*

$$\lambda_k(v) = \begin{cases} 0 & v \in [\underline{v}, \tilde{v}_{k-1}) \cup (\tilde{v}_k, \bar{v}] \\ \frac{b(v)^\varepsilon a_k(v)^{\varepsilon-1}}{B_k(\tilde{v}_{k-1}, \tilde{v}_k)^\varepsilon} \pi_k & v \in (\tilde{v}_{k-1}, \tilde{v}_k), \end{cases}$$

²¹As before, we constrain the set of mechanisms available to the government to ones that deterministically condition upon worker incomes. This assumption is standard in the literature and to a first approximation describes current tax codes. In our setting, it implies that the government cannot observe the task a worker does or the amount of task output. The former may reasonably reflect the inherent difficulties in distinguishing between a worker's formal job description and the tasks that the worker actually performs.

where $B_k(\tilde{v}_{k-1}, \tilde{v}_k) := \left[\int_{\tilde{v}_{k-1}}^{\tilde{v}_k} b(v)^\varepsilon a_k(v)^{\varepsilon-1} dv \right]^{\frac{1}{\varepsilon}}$. All workers of talent k earn a common wage $w_k = \omega(v) a_k(v)$, $v \in [\tilde{v}_{k-1}, \tilde{v}_k]$. Relative wages are given by

$$(14) \quad \frac{w_{k+1}}{w_k} = \frac{a_{k+1}(\tilde{v}_k)}{a_k(\tilde{v}_k)} = \frac{B_{k+1}(\tilde{v}_k, \tilde{v}_{k+1}) / \{\pi_{k+1} e_{k+1}\}^{\frac{1}{\varepsilon}}}{B_k(\tilde{v}_{k-1}, \tilde{v}_k) / \{\pi_k e_k\}^{\frac{1}{\varepsilon}}}.$$

Conditional on the effort profile $\{e_k\}$, the equilibrium allocation of talent to tasks maximizes output.

PROOF:

See online Appendix I.

Efficiency of assignment (in the sense of output maximization) conditional on effort implies that output is given by the following indirect production function over efforts:

$$(15) \quad F(\pi_1 e_1, \dots, \pi_K e_K) = \sup \left\{ A \left\{ \sum_{k=1}^K B_k(\tilde{v}_{k-1}, \tilde{v}_k) \{e_k \pi_k\}^{\frac{\varepsilon-1}{\varepsilon}} \right\}^{\frac{\varepsilon}{\varepsilon-1}} \mid \text{s.t. } \underline{v} \leq \tilde{v}_1 \leq \dots \leq \tilde{v}_{K-1} \leq \bar{v} \right\}.$$

With F determined in this way, the environment effectively reduces to that in Section II and the government's problem to (PP). Recovery of an optimal tax equilibrium can be decomposed into two steps. The outer step is simply (PP) at the induced production function F ; the embedded inner step solves the assignment problem (15) at each candidate effort allocation $\{e_k\}$ and, hence, evaluates F at $\{e_k\}$.²²

In contrast to Section II, the production function F is micro-founded; changes in parameters of this production function can be related to changes in the demand for tasks b and the productivity of task-talent matches $\{a_k\}$. The inner step assignment problem is essentially the same as those considered in Teulings (1995); Costinot and Vogel (2010); and Acemoğlu and Autor (2011) (with the distinction that the supply of each talent's labor is selected as part of an optimal tax equilibrium rather than being pinned down parametrically).²³ Solving the assignment problem at an effort profile $\{e_k\}$ reduces to finding a sequence of task thresholds $\{\tilde{v}_k\}_{k=1}^{K-1}$ satisfying the discrete boundary value problem:

$$(16) \quad \frac{a_{k+1}(\tilde{v}_k)}{a_k(\tilde{v}_k)} = \frac{B_{k+1}(\tilde{v}_k, \tilde{v}_{k+1}) / \{\pi_{k+1} e_{k+1}\}^{\frac{1}{\varepsilon}}}{B_k(\tilde{v}_{k-1}, \tilde{v}_k) / \{\pi_k e_k\}^{\frac{1}{\varepsilon}}},$$

with $\tilde{v}_0 = \underline{v}$ and $\tilde{v}_K = \bar{v}$.

²²In a tax equilibrium, a worker reproducing the income and paying the taxes of a less talented type will exert less effort in the task that pays her the best wage, she does not move to the task of the less talented whose income she mimics. Thus, worker task (and wage) choice is independent of the effort she exerts and the income she earns in the task. The counterpart of this in the decomposition just described is the incentive constraint in the outer step which depends on relative wages and only via them on task choice.

²³In fact the analysis on p. 758–60 of Costinot and Vogel (2010) in which the labor input across “skills” is changed in particular ways represents a partial exploration of the indirect production function.

An immediate consequence of Proposition 2 and the absolute advantage condition Assumption 1(ii) is that $\frac{w_{k+1}}{w_k} = \frac{a_{k+1}(\tilde{v}_k)}{a_k(\tilde{v}_k)} > 1$. Consequently, talents are strictly ordered by equilibrium wages and “wage pooling” (the payment of the same wage to different talent types) does not occur.²⁴ Proposition 1 identifies relative wage-effort elasticities $\phi_{k,j}$ as key determinants of the wage compression channel and, hence, marginal taxes. If each $\log(a_{j+1}/a_j)$ is differentiable, then in a tax equilibrium the terms $\phi_{k,j}$ can be expressed as

$$(17) \quad \phi_{k,j} = -\frac{\partial \log(w_{j+1}/w_j)}{\partial \log e_k} = \begin{cases} -\frac{\partial \log(a_{j+1}/a_j)}{\partial \log \tilde{v}_j} \prod_{l=k}^{j-1} \left(\frac{\partial \log \tilde{v}_{l+1}}{\partial \log \tilde{v}_l} \right) \frac{\partial \log \tilde{v}_k}{\partial \log e_k} & j \geq k \\ -\frac{\partial \log(a_{j+1}/a_j)}{\partial \log \tilde{v}_j} \prod_{l=j}^{k-2} \left(\frac{\partial \log \tilde{v}_l}{\partial \log \tilde{v}_{l+1}} \right) \frac{\partial \log \tilde{v}_{k-1}}{\partial \log e_k} & j < k \end{cases}.$$

Thus, elasticity $\phi_{k,j}$ depends upon the local comparative advantage of talents j and $j+1$ at the threshold \tilde{v}_j , the sensitivity of the $k-1$ th or k th task threshold to the effort of the k th talent $\frac{\partial \log \tilde{v}_k}{\partial \log e_k}$ and the sensitivity of thresholds intermediate between j and k to one another $\frac{\partial \log \tilde{v}_{l+1}}{\partial \log \tilde{v}_l}$.

Only under very special conditions is the induced production function F a CES function. One such case occurs when each $\frac{\partial \log(a_{j+1}/a_j)}{\partial \log \tilde{v}_j} = 0$, $\phi_{k,j} = 0$, talents are perfectly substitutable across tasks and F is linear. Another²⁵ occurs when the a_k functions are indicators for the sub-intervals $[\underline{v}, \tilde{v}_1]$, $(\tilde{v}_1, \tilde{v}_2]$, \dots , $(\tilde{v}_{K-1}, \bar{v}]$. Then workers are as substitutable as the tasks into which they are locked. For more general cases, however, relative wage-effort elasticities are functions of technological parameters and the effort profile $\{e_k\}$ and, hence, indirectly policy. Thus, they are not structural. In the Appendix, we prove:

LEMMA 1: *Each $\frac{\partial \log \tilde{v}_j}{\partial \log \tilde{v}_{j+1}}$, $\frac{\partial \log \tilde{v}_{j+1}}{\partial \log \tilde{v}_j}$ and $\frac{\partial \log \tilde{v}_k}{\partial \log e_k}$ is positive. Each $\frac{\partial \log \tilde{v}_{k-1}}{\partial \log e_k}$ is negative. If $\frac{\partial \log(a_{j+1}/a_j)}{\partial \log \tilde{v}_j} > 0$, then $\phi_{k,j} < 0$ if $j \geq k$ and $\phi_{k,j} > 0$ if $j < k$. In addition, $\phi_{k,k} \in [-1/\varepsilon, 0]$ and $\phi_{k,k-1} = [0, 1/\varepsilon]$.*

PROOF:

See online Appendix I.

The economics behind Lemma 1 is straightforward. Consider a small increase in e_k (perhaps in response to a policy change). This raises output in tasks $[\tilde{v}_{k-1}, \tilde{v}_k]$, placing downward pressure on $[\tilde{v}_{k-1}, \tilde{v}_k]$ -shadow prices and, hence, the wage w_k of

²⁴ Assumption 1(ii) (i.e., global absolute advantage of more talented types across the entire task space) is sufficient, but not necessary for this result. Local absolute advantage of successive talents $k+1$ at each task boundary \tilde{v}_k is enough.

²⁵ Although, this case is not consistent with talent-complexity comparative advantage (except when $K = 2$), smoothness or continuity of the a_k functions.

talent k workers. These workers respond by populating tasks that are both below \tilde{v}_{k-1} and above \tilde{v}_k . This task migration moderates, but does not fully offset the impact of the increase in e_k on w_k . As k -talents move into less complex tasks in which they have a comparative disadvantage relative to $k - 1$ -talents and more complex tasks in which they have a comparative disadvantage relative to $k + 1$ -talents so w_k/w_{k-1} falls and w_{k+1}/w_k rises. Moreover, as k -talents spill into neighboring tasks, output of these tasks increases, depressing their shadow prices and inducing neighboring talents to migrate into new tasks. Workers of talent $k + 1$ move into tasks above \tilde{v}_{k+1} , while workers of talent $k - 1$ talents move into tasks below \tilde{v}_{k-1} . A ripple effect is created with each task threshold \tilde{v}_j above k rising and each threshold below k falling. Since relative wages between adjacent talents are determined by productivity ratios at thresholds (i.e., by $a_{j+1}(\tilde{v}_j)/a_j(\tilde{v}_j)$), an effort change by talent k workers can affect relative wages across the whole spectrum of talents and be a motive for encouraging or discouraging that talent's effort.

Expressions for the threshold elasticities $\frac{\partial \log \tilde{v}_j}{\partial \log \tilde{v}_{j+1}}$, $\frac{\partial \log \tilde{v}_{j+1}}{\partial \log \tilde{v}_j}$, and $\frac{\partial \log \tilde{v}_k}{\partial \log e_k}$ are given in the proof of Lemma 1. They point to the role of the parameters b and a in influencing the sensitivity of task choices and, hence, relative wages to a given talent's effort. Suppose that workers of talent j migrate into more complex tasks either because they have increased their effort or because the tasks that they originally performed have been encroached upon by $j - 1$ talents. If there is much demand and, hence, high b -values for tasks immediately above \tilde{v}_j , then these tasks will soak up this migration with little change in the threshold \tilde{v}_j . Conversely, if b -values in this neighborhood are low, then talent j -workers will migrate further up through the task set pushing \tilde{v}_j to a new possibly much higher level. In the former case, the impact on the w_{j+1}/w_j wage differential will be muted; in the latter case, it will be enhanced. Turning to the a function, an increase in the comparative advantage of talent in complex tasks, raises $\frac{\partial \log(a_{j+1}/a_j)}{\partial \log \tilde{v}_j}$ and, hence, the sensitivity of relative wages to task threshold adjustment. The resulting upwards pressure on $\phi_{k,j}$ is dampened by the deterrence to task migration and task threshold adjustment and, hence, lower values for $\frac{\partial \log \tilde{v}_{j+1}}{\partial \log \tilde{v}_j}$ and $\frac{\partial \log \tilde{v}_k}{\partial \log e_k}$, created by higher comparative advantage.

C. An Example: Technical Change in the Two Talent Model

We now make some of the preceding observations more precise in the context of a simple two talent example. For concreteness, we label these talents low ($k = L$) and high ($k = H$) rather than 1 and 2. We restrict preferences to be quasi-linear in consumption, $U(c, e) = c - \frac{e^{1+\gamma}}{1+\gamma}$, with $\gamma > 0$ and denote the government's Pareto weights by g_k , $k \in \{L, H\}$. To create a motive for redistribution to low skills, we assume $g_L > \pi_L$. In this case, the Mirrlees and Wage Compression components can be consolidated to give the (Stiglitz) optimal tax functions:

$$(18) \quad \frac{\tau_L^*}{1 - \tau_L^*} = \left(\frac{g_L}{\pi_L} - 1 \right) \left\{ 1 - \left(\frac{1}{\mathcal{W}^*} \right)^{1+\gamma} \left\{ 1 - \frac{1}{\mathcal{E}^*} \right\} \right\} \geq 0,$$

$$(19) \quad \frac{\tau_H^*}{1 - \tau_H^*} = \left(\frac{g_H}{\pi_H} - 1 \right) \left(\frac{1}{\mathcal{W}^*} \frac{e_L^*}{e_H^*} \right)^{1+\gamma} \frac{1}{\mathcal{E}^*} \leq 0,$$

where $\mathcal{W}^* = \frac{w_H^*}{w_L^*}$ is the optimal talent premium and the substitutability of talents at the optimum is completely described by the elasticity of substitution \mathcal{E}^* . In the assignment setting, both \mathcal{W}^* and \mathcal{E}^* are endogenous. If Assumption 1 is maintained, then in an optimal tax equilibrium the set of tasks is partitioned at a threshold \tilde{v}^* , with low talents working in tasks below \tilde{v}^* and high talents working in tasks above. The talent premium satisfies $\mathcal{W}^* = \frac{a_H(\tilde{v}^*)}{a_L(\tilde{v}^*)}$, while the elasticity of substitution between the labor of the two talents is given by $\mathcal{E}^* = \mathcal{E}(\tilde{v}^*, a, b)$, where

$$(20) \quad \mathcal{E}(\tilde{v}; a, b) := -\frac{\partial \log \frac{e_H}{e_L}}{\partial \log \frac{w_H}{w_L}} = \varepsilon + \frac{1}{\frac{\partial \log a_H/a_L}{\partial v}(\tilde{v})} \left[\frac{b_H(\tilde{v})}{B_H(\tilde{v})} + \frac{b_L(\tilde{v})}{B_L(\tilde{v})} \right] \geq \varepsilon,$$

$B_L(\tilde{v}) := \int_{\underline{v}}^{\tilde{v}} b(v)^\varepsilon a_L(v)^{\varepsilon-1} dv$, $B_H(\tilde{v}) := \int_{\tilde{v}}^{\bar{v}} b(v)^\varepsilon a_H(v)^{\varepsilon-1} dv$ and for $k \in \{L, H\}$, $b_k(\tilde{v}) := b(\tilde{v})^\varepsilon a_k(\tilde{v})^{\varepsilon-1}$. Equation (20) makes explicit the role of task migration in raising this elasticity of substitution above ε , the elasticity of substitution between task outputs. Specifically, the term $\frac{\partial \log (a_H/a_L)}{\partial v}(\tilde{v})$ is the local comparative advantage of high talents in the neighborhood of the threshold task \tilde{v} . If this term equals ∞ , then workers are as substitutable as the interval of tasks into which they are locked. Otherwise, their ability to migrate across tasks enhances their substitutability. Equation (20) highlights the dependence of the elasticity of substitution on technological parameters and its (implicit) dependence on policy.

The workers' equilibrium first order conditions in this setting together with (14) gives

$$(21) \quad \mathcal{W}^* = \frac{a_H(\tilde{v}^*)}{a_L(\tilde{v}^*)} = \left(\frac{B_H(\tilde{v}^*)}{B_L(\tilde{v}^*)} \right)^{\frac{\gamma\varepsilon}{1+\gamma\varepsilon}} \left(\frac{\pi_L}{\pi_H} \right)^{\frac{\gamma}{1+\gamma\varepsilon}} \left(\frac{1 - \tau_H}{1 - \tau_L} \right)^{\frac{1}{1+\gamma\varepsilon}}.$$

Equation (21) gives the threshold \tilde{v}^* and relative wages as functions of a and b and relative marginal taxes. Substituting for optimal marginal taxes from (18) and (19) reduces the system of equilibrium equations to a single equation in one unknown, \tilde{v}^* :

$$(22) \quad \frac{a_H(\tilde{v}^*)}{a_L(\tilde{v}^*)} = \left(\frac{B_H(\tilde{v}^*)}{B_L(\tilde{v}^*)} \right)^{\frac{\gamma\varepsilon}{1+\gamma\varepsilon}} \left(\frac{\pi_L}{\pi_H} \right)^{\frac{1+\gamma}{1+\gamma\varepsilon}} \times \left\{ \frac{g_H - (\pi_H - g_H) \left(\frac{a_H(\tilde{v}^*)}{a_L(\tilde{v}^*)} \right)^{(1+\gamma)(\varepsilon-1)} \left(\frac{B_H(\tilde{v}^*)}{B_L(\tilde{v}^*)} \right)^{-(1+\gamma)\varepsilon} \frac{1}{\mathcal{E}^*}}{g_L - (g_L - \pi_L) \left(\frac{a_L(\tilde{v}^*)}{a_H(\tilde{v}^*)} \right)^{1+\gamma} \left(\frac{\mathcal{E}^* - 1}{\mathcal{E}^*} \right)} \right\}^{\frac{1}{1+\gamma\varepsilon}}.$$

It follows easily from (22) that if $\varepsilon \geq 1$ (so that goods, and, hence, efforts of different talents are gross substitutes) and if $\mathcal{E}(\cdot; a, \cdot)$ is (locally) constant, then complexity-biased perturbations of b that raise B_H/B_L lead to increases in \tilde{v}^* and \mathcal{W}^* . Intuitively, increases in the relative demand for more complex tasks raise the relative shadow price of such tasks and encourage less talented workers to migrate into them (\tilde{v}^* rises). However, such task-upgrading erodes the comparative advantage of low talents and the talent premium (\mathcal{W}^*) rises. These effects are mitigated by adjustments in relative efforts that occur in response to wage adjustments and that are reinforced by changes to tax policy. Overall, a rise in $B_H/B_L(\cdot)$ is associated with a higher talent premium, a tightening of the incentive constraint between low and high talents and higher marginal taxes on low talents.

Sufficient conditions for the elasticity of substitution \mathcal{E}^* to be constant in response to a shift in task demand are rather stringent.²⁶ In general, it may rise or fall as a direct effect of the change in b or the indirect effect of changes in \tilde{v}^* on

$\frac{1}{\frac{\partial \log a_H/a_L}{\partial v}(\tilde{v}^*)} \left\{ \frac{b_H(\tilde{v}^*)}{B_H(\tilde{v}^*)} + \frac{b_L(\tilde{v}^*)}{B_L(\tilde{v}^*)} \right\}$ in (20). These changes may reinforce or offset the responses just described. To the extent that \mathcal{E}^* is increased, the government's ability to compress wage differentials and relax incentive constraints is reduced. It is correspondingly encouraged to reduce relative taxation of low talents and to permit a further increase in the talent premium. The reverse is true if \mathcal{E}^* falls.

Turning next to the consequences of variation in a , suppose that $\log \frac{a_H(v)}{a_L(v)} = \alpha_1 + \alpha_2(v - \underline{v})$ so that α_1 controls the absolute advantage of high talents (in the lowest task) and α_2 controls their comparative advantage in more complex tasks. If $\mathcal{E} > 1$ and $\left\{ \frac{b_H(\tilde{v}^*)}{B_H(\tilde{v}^*)} + \frac{b_L(\tilde{v}^*)}{B_L(\tilde{v}^*)} \right\}$ is locally constant, then small technologically induced increases in α_2 will, from (22), both raise the talent premium and reduce the elasticity of substitution \mathcal{E}^* .²⁷ Low-talent marginal taxes τ_L^* will rise both because \mathcal{W}^* rises and because the wage compression channel is enhanced via the reduction in \mathcal{E}^* : as workers become less substitutable, the government is encouraged to offset the rise in the talent premium by discouraging low-talent effort through taxation. Increases in α_1 work in a related way, but absent any reinforcing adjustment in \mathcal{E}^* . As in the case of complexity-biased perturbations in the b functions, adjustments in the $\frac{b_H(\tilde{v}^*)}{B_H(\tilde{v}^*)} + \frac{b_L(\tilde{v}^*)}{B_L(\tilde{v}^*)}$ term (either direct through changes to the b_k functions or indirect through adjustments to \tilde{v}^*) may work to reinforce or dampen these effects.

Summary.—Technical change that increases the talent wage premium and reduces the substitutability of talents is associated with higher optimal marginal taxes on low talents. Change that increases both the talent income premium and the substitutability of talents is associated with lower marginal subsidies on high talents. In general,

²⁶For example, if $\varepsilon = 1$, b is constant and equal to one and the remaining parameters are such that $\tilde{v}^* = 1/2$, then \mathcal{E}^* is locally constant.

²⁷In this case the task threshold \tilde{v}^* falls: the increased productivity of high talents in complex tasks reduces the relative shadow price of such tasks and encourages high talents to downgrade their tasks. Despite some erosion of their comparative advantage, their relative wages rise.

the technical parameters a and b influence both talent premia and talent substitutability directly and indirectly through endogenous task assignment. The analysis is more complicated in settings with multiple talents. Such settings are, however, essential for exploring the policy implications of recently documented polarizing shifts in the pattern of wages and employment across occupations.

D. Comparison to Rothschild and Scheuer (2013)

We briefly describe the connections between our model and that of Rothschild and Scheuer (2013). Our model features a continuum of tasks and a finite set of talents, but it is readily reformulated as one with a continuum of tasks and talents (see online Appendix II). In both formulations our assumptions ensure that the ordering over talents translates directly into an ordering over wages. Consequently, the pattern of (local) binding incentive constraints over talents is easily inferred and consumption and effort allocations can be solved directly as functions of talent. We use our approach to relate optimal taxes to the indirect production function F , relative wage elasticities and, hence, properties of the task-talent productivity function a .

In contrast, Rothschild and Scheuer (2013) consider an environment with a finite number of tasks in which an agent's type is her productivity in each task and is, thus, multidimensional. In this case, the structure of binding incentive-compatibility conditions across allocations expressed as functions of type is complicated. However, such conditions become quite standard if consumption and effort allocations are reformulated as functions of wages. The cost of this reformulation is a rather complicated joint restriction on allocations and the (endogenous) distribution over wages. To solve such a problem Rothschild and Scheuer (2013) propose a quite different inner-outer method than that used here. In the inner step the allocation of labor across tasks and, hence, the wage distribution is fixed and an optimal incentive-compatible consumption-effort profile (over wages) consistent with this allocation is found. In the outer step, the labor allocation and the wage distribution are determined. Rothschild and Scheuer (2013) use this approach to relate optimal taxes to the impact of effort on the wage distribution. Their formula, thus, provides an alternative perspective on the forces shaping tax policy in an endogenous wage environment to ours.

Rothschild and Scheuer's (2013) model permits intra-task wage dispersion and task-specific wage distributions with overlapping support. It thus allows the implications of these things for policy to be explored, ours does not. On the other hand, our model connects directly to the technical change literature in labor economics. It underpins an empirical strategy for quantifying the effect of technical change on optimal policy described in Sections IV and V. Thus, the models are complementary. In online Appendix VI, we present a general formulation that nests our model and that of Rothschild and Scheuer (2013) and makes transparent the alternative approaches taken. It then specializes that formulation to one intermediate between our model and theirs. This formulation incorporates intra-task wage dispersion and wage overlap, while preserving our approach to formalizing the impact of technical change.

IV. Measuring Technical Change

In this section, we measure the extent of technical change in the United States. Our data source is the Current Population Survey (CPS).²⁸ We proceed as if this data was generated by a (possibly suboptimal) tax equilibrium and use parametric assumptions and equilibrium restrictions from our model to identify and estimate the technological parameters a and b in the 1970s and 2000s. In Section V, we calculate optimal tax equilibria at these estimated parameters.

A. Determining Types and Tasks

Mapping Empirical Occupations to Ordered Sets of Tasks.—The CPS categorizes workers into distinct occupations; our sample contains $M = 302$ occupations. The CPS also provides information on worker earnings and hours worked from which a measure of wages can be imputed. Our model involves an interval of tasks ordered by complexity. We identify tasks with empirical occupations and use the average wage paid in each occupation to infer its complexity. In so doing, we utilize the model's implication that task wages are rising in task complexity. We normalize the task space to $[\underline{v}, \bar{v}] = [0, 1]$ and subdivide this interval into M subintervals of length $\Delta v = \frac{1}{M}$, $\mathcal{V}_m = [v_{m-1}, v_m]$. We calculate the imputed average wage in each occupation using 1970's data and rank occupations according to this wage. The m th ranked occupation is then mapped to the m th subinterval \mathcal{V}_m .²⁹

We use data on the skill content of occupations contained in the O*NET database to corroborate our inferred complexity ordering over occupations. The O*NET database provides a detailed description of the skill (35 distinct skills are considered) and ability (52 distinct abilities are considered) content of each occupation.³⁰ We recover from the O*NET a single index describing the importance of each skill and ability for each occupation.³¹ We then calculate the correlations of these skill/ability indices with our wage imputed rank. We find that the three most correlated skills (correlation in parenthesis) are: *complex problem solving* (0.66); *critical thinking* (0.62); and *judgment and decision making* (0.61). The three most correlated abilities are: *deductive reasoning* (0.63); *inductive reasoning* (0.60); and *written comprehension* (0.57). The least correlated skill is *equipment maintenance* (−0.07), while the least correlated abilities are: *stamina* (−0.33) and *trunk strength* (−0.37). These correlations suggest that the average wage paid in an occupation is informative about that occupation's complexity.

²⁸ King et al. (2010). Further details of our use and treatment of the data are given in online Appendix IIIA.

²⁹ We keep this ranking over occupations fixed. In doing so, we follow the precedent of Acemoglu and Autor (2011). Fixing the ranking allows us to unambiguously identify an index v with a physical occupation and to interpret variations in the parameters a and b as occurring in a given physical occupation rather than at a given complexity index whose physical interpretation is shifting. However, there is some reranking of occupations over time in the data. In online Appendix IIIC we describe the implications of using current rather than the 1970s wage ranking for our estimates of the a and b functions and for optimal taxes.

³⁰ The O*NET database contains 974 occupations. We relate these to the occupations contained in CPS in two steps. We first map the occupations in our sample to the Standard Occupation Classification of the 2000 census. We then map these occupations to those in the nineteenth release of the O*NET. A small number of occupations are recoded manually. We thank Giovanni Gallipoli for directing us toward the O*NET.

³¹ Specifically, the index is the product of the importance and level measures in O*NET.

Recovering the Empirical Assignment Function \tilde{v} .—The model in Section III featured a finite number of talents; this facilitated the derivation of analytical results. However, for the remainder of the paper we find it convenient to treat worker talent symmetrically with task complexity and to assume that workers are distributed uniformly across an interval of talents, $k \in [\underline{k}, \bar{k}]$.³² Thus, a worker's talent should now be interpreted as an index (and a rank), the implications of which for productivity are captured by the function: $a: [\underline{k}, \bar{k}] \times [\underline{v}, \bar{v}] \rightarrow \mathbb{R}_+$. Although the distribution over the (ordinal) talent index is uniform, the distribution over (cardinal) productivities is not: it is induced endogenously by a and by the assignment of talent to tasks. The set $[\underline{k}, \bar{k}]$ is normalized to $[0, 1]$.

The continuous analogue of the task thresholds $\{\tilde{v}_k\}$ is a task assignment function $\tilde{v}: [\underline{k}, \bar{k}] \rightarrow [\underline{v}, \bar{v}]$. This function is strictly increasing in our model. Denote its inverse by \tilde{k} . Under the assumption that workers are distributed uniformly across talent indices, \tilde{k} is the distribution of workers across tasks. Consequently, we treat the distribution of workers across ordered occupations as the empirical counterpart of \tilde{k} and \tilde{v} to be the inverse of this.

B. Estimating b

It is well known that the elasticity of substitution between goods and factor augmenting technical progress (ε and b in our case) cannot be separately identified from data on outputs, inputs, and marginal products—an observation that goes back to Diamond, McFadden, and Rodriguez (1978). In our baseline case, we restrict the elasticity of substitution between task outputs, ε , to be one (so that the final good production function is Cobb-Douglas) and identify $b(v)$ with the share of total compensation paid to workers in task v .³³ Thus, estimates of b may be calculated from compensation shares independently of knowledge of the a s. Specifically, under the Cobb-Douglas restriction, the firm's first order conditions from the continuous-talent version of (10), imply for almost all (k, v) :


$$(23) \quad \omega(k, v) = Y \frac{a(k, v)b(v)}{y(v)}.$$

In the continuous talent setting, task output is given by $y(v) = a(\tilde{k}(v), v)e(\tilde{k}(v))\tilde{k}_v(v)$, with \tilde{k}_v the derivative of \tilde{k} . Combining this with (23) and integrating over \mathcal{V}_m gives total labor income in occupation m in terms of the b -function:

$$\int_{\mathcal{V}_m} \omega(\tilde{k}(v), v)e(\tilde{k}(v))\tilde{k}_v(v) dv = Y \int_{\mathcal{V}_m} b(v) dv.$$

³²The convenience is two-fold. First, since occupational (task) data is discrete, assuming a continuous set of talents avoids having to deal with talent groups that are distributed across adjacent occupations. Second, it allows us to apply numerical optimal control methods to solve the problem. A formal statement of the continuous talent-continuous task model can be found in online Appendix II.

³³The quantitative implications of alternative assumptions for ε are considered in online Appendix IIB.

FIGURE 3. EVOLUTION OF $\log(b(v))$ ACROSS DECADES

Average income in occupation m , i_m , is then obtained by dividing both sides by the mass of workers in the occupation, S_m :

$$i_m := \frac{1}{S_m} \int_{\mathcal{V}_m} \omega(\tilde{k}(v), v) e(\tilde{k}(v), v) \tilde{k}_v(v) dv = \frac{Y}{S_m} \int_{\mathcal{V}_m} b(v) dv.$$

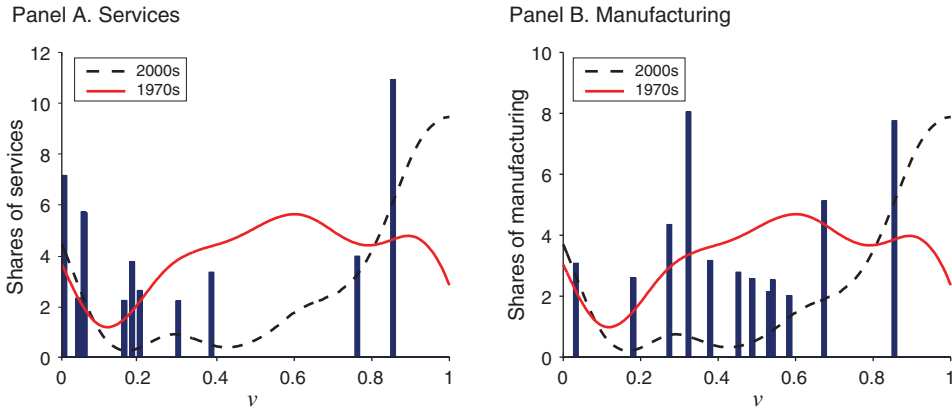
Thus, the average value of b in occupation m , $b_m \Delta v := \int_{v_{m-1}}^{v_m} b(v) dv$, is

$$(24) \quad b_m = \frac{S_m i_m}{\Delta v Y}, \quad \forall m = 1, \dots, M.$$

We identify Y with per capita labor income.³⁴ A smooth estimate of the b -function is obtained by fitting a LOWESS model to $\{v_m, \log b_m\}$ data.³⁵ Figure 3 displays estimates of b for the 1970s and the 2000s. The figure shows that b rises (slightly) for low and (significantly) for high v -occupations, but falls for intermediate ones.

³⁴In 2005 dollars we have $Y_{70} = \$36,998$ and $Y_{00} = \$45,260$. M is 302. In aggregate data using GDP deflator (Table 1.1.9 in NIPA) and total nonfarm payroll (BLS) we get a value of real compensation per worker equal to $Y_{70} = \$37,114$ and $Y_{00} = \$53,304$. However deflating using CPI we get values consistent with our sample: $Y_{70} = \$37,966$ and $Y_{00} = \$45,151$.

³⁵The LOWESS scatterplot smoothing builds up a smooth curve through a set of data points by fitting simple linear or quadratic models to localized subsets of data. We use a smoothing parameter of 0.4.

FIGURE 4. OCCUPATIONS AND ν

Notes: Histograms: shares of occupations over ν . Plots: smoothed values for $\log(b(\nu))$ over ν and across decades.

The picture is consistent with the phenomenon of *job polarization* as discussed in Section I. This polarization feature is robust to different sample selection assumptions, see online Appendix IIIA for details.

Figure 4 sharpens intuition concerning the relation of different ν s to the data. The figure overlays the values of $b(\cdot)$ with a bar graph displaying the employment shares of occupations belonging to particular sectors. Figure 4, panel A does this for services and Figure 4, panel B for manufacturing.³⁶ The service sector is associated mostly with extreme and, especially, “low” ν occupations (the bar on the right in Figure 4, panel A refers to managers and administrative support), while manufacturing is mostly middle ν occupations (although with a wider range).

C. Estimating a

The envelope condition from the task choice component of the worker’s equilibrium problem, $w(k) = \max_{\nu \in [\underline{\nu}, \bar{\nu}]} \omega(k, \nu)$, implies that

$$(25) \quad \frac{d \log w}{dk}(k) = \frac{\partial \log \omega(k, \tilde{\nu}(k))}{\partial k} = \frac{\partial \log a(k, \tilde{\nu}(k))}{\partial k} = \frac{\partial \alpha}{\partial k}(k, \tilde{\nu}(k)),$$

where $\alpha(k, \nu) := \log a(k, \nu)$. An empirical counterpart for $\frac{d \log w}{dk}$ is constructed in three steps. First, information from the CPS on weeks and usual hours worked in the previous year and self reported yearly labor income is used to impute workers’ average hourly wages. Second, wages are averaged over occupation to construct empirical counterparts of $w(\tilde{k}(\nu))$. Third, a LOWESS smoother is applied to the log of this series and to \tilde{k} , derivatives of each function are calculated

³⁶Not shown are occupations that constitute less than 2 percent of the workforce of each sector.




FIGURE 5. LOWESS SMOOTHED LOG (WAGES) OVER TALENTS

and $\frac{d \log w(\tilde{k}(v))}{dk} = \frac{d \log w(\tilde{k}(v))}{dv} \bigg/ \frac{\partial \tilde{k}(v)}{\partial v}$ is found. Figure 5 displays the empirical values for (smoothed) $\log w(k)$ for the 1970s and the 2000s. From the twentieth to the eightieth talent percentile, this function is roughly linear in k in the 1970s and remains so in the 2000s. In the 1970s, it steepens over the top talent decile, while in the 2000s, it steepens over the top *two* deciles. In addition, for both decades, but especially for the 1970s, the profile is steeper over the bottom two deciles.

The evolution of $\log w(k)$ shown in Figure 5 suggests that between the 1970s and the 2000s the wages of low-ranked talents caught up with mid-ranked talents, while the wages of mid-ranked talents fell behind those at the top. These developments are qualitatively consistent with a fall in the returns to talent in simpler tasks combined with an increase in talent-complexity comparative advantage (so that talent premia rise in the most complicated tasks and occupations). This motivates us to select:

$$(26) \quad \frac{\partial \alpha}{\partial k}(k, v) = \alpha_1 + \alpha_2 \cdot v.$$

Here, α_1 captures the return to pure talent, while α_2 captures comparative advantage.³⁷ We recover estimates of α_1 and α_2 by regressing $\frac{d \log w(\tilde{k}(v))}{dk}$ onto a constant and the task index v . The regression is weighted by the share of workers in each v .

³⁷ Online Appendix IIIE considers a case in which comparative advantage is increasing with task complexity.

TABLE 1—ESTIMATION OF PRODUCTIVITY FUNCTION

	α_1	α_2
1970s	1.07 (0.25)	1.71 (0.28)
2000s	0.42 (0.32)	3.01 (0.22)

Notes: $N = 302$. Estimation of α_1 and α_2 from (25). Standard errors in parenthesis.

Source: Authors' calculations.

Results are reported in Table 1. They show a significant increase in the comparative advantage parameter α_2 between the 1970s and 2000s. Loosely, this is driven by the increase in wage growth over high talents occurring between the 1970s and the 2000s.³⁸ In online Appendix IIID, we look outside of the CPS for corroborating evidence of increasing comparative advantage. Specifically, we use data on the change in the skill/ability content of occupations contained in different editions of the O*NET database. We find evidence that the use and importance of skills and abilities associated with complex tasks has increased in high-wage occupations relative to low.

Finally, the parameter A is given by the ratio of per capita income to the approximation of the CES aggregator $\exp \left\{ \int_{\underline{v}}^{\bar{v}} b(v) \log \{y(v)\} dv \right\}$.

V. Quantitative Implications for Policy

In this section, we compute optimal policy responses to the technical change estimates derived in Section IV. Calculation of policy requires a specification of worker and societal preferences and the amount of resources devoted to public spending. We briefly turn to this and then give our quantitative results.

A. Selection of Other Parameters and Computational Method

We assume that worker preferences are given by: $U(c, e) = \log c - \frac{e^{1+\gamma}}{1+\gamma}$. Note that the choice of U has no impact on the estimation of $b(v)$ and $a(k, v)$. We follow Chetty et al. (2011) and set the Frisch labor supply elasticity to $1/\gamma = 0.75$. We identify the share of output allocated to public spending with the aggregate tax to income ratio in our CPS sample. On this basis, $(G/Y)_{70} = 16.2$ percent and $(G/Y)_{00} = 14.0$ percent; we set the G/Y ratio to the intermediate value of 15 percent.³⁹ Finally, in our benchmark calculations a utilitarian government is assumed: $g_k = \pi_k$ for all talents k .

³⁸ Kaplan and Rauh (2013) emphasize the rise of “superstar” pay across a variety of high-income occupations. In our empirical strategy “superstar” workers belong to (measured) occupations inhabited by much lower paid workers. It is arguable that these different workers trade in distinct task-markets with distinct shadow prices. The implication of this is a downward bias in the estimate of comparative advantage (α_2). Given the evolution of inequality in the United States, this bias is likely to be more significant for the 2000s.

³⁹ NIPA data (Table 1.1.6) gives $(G/Y)_{70} = 23.9$ percent and $(G/Y)_{00} = 19.3$ percent. However, since we are concerned with spending financed out of income taxation (paid by our subsample of labor income earners) we use the alternative CPS-generated estimates.

TABLE 2—OPTIMAL TAX RATES IN REAL LABOR INCOME

	Decade	Percentiles of income					
		10th	25th	50th	75th	90th	99th
Averages	1970s	−11.9	−7.3	6.9	22.3	26.1	22.3
	2000s	−2.3	−1.1	5.6	19.9	26.1	21.9
Marginals	1970s	20.3	34.1	44.3	40.3	23.9	−0.6
	2000s	15.3	25.4	39.7	42.2	27.4	−2.2

Source: Authors' calculations.

To calculate optimal policy at our selected and estimated parameters, we first formulate the government's optimization as an optimal control problem. Details of this formulation are given in online Appendix II. We then solve the problem numerically using the GPOPS-II software.⁴⁰

B. Optimal Tax Results

Table 2 reports optimal average and marginal tax rates as a function of income percentiles for the 1970s and the 2000s. Over this time period, average rates rise at low incomes and fall at high and, especially, middle incomes. Transfers to the lowest deciles are reduced. Overall, the reform favors those in the middle. Marginal rates fall at low to mid incomes and rise at higher incomes. In the extreme tails they move in the opposite directions: rising in the very lowest and falling in the very highest (where marginal subsidies are increased) percentiles. To understand the evolution of optimal tax reported in Table 2, we return to the tax formula (4) derived earlier.

Accounting for Optimal Taxes.—Tax formula (4) allows us to decompose optimal tax rates into “Mirrleesian” and “Wage Compression” components. In particular, let τ_k^M denote the “Mirrleesian” marginal tax rate in the absence of the wage compression term:⁴¹

$$\tau_k^M = \frac{\frac{\Delta w_{k+1}^*}{w_{k+1}^*} \frac{1 - \Pi_k}{\pi_k} \mathcal{H}_k^* \Psi_k^*}{1 + \frac{\Delta w_{k+1}^*}{w_{k+1}^*} \frac{1 - \Pi_k}{\pi_k} \mathcal{H}_k^* \Psi_k^*}.$$

The tax rate τ_k^M is that which an optimizing government would apply if wages were fixed at their optimal levels $\{w_k^*\}$. Define the wage compression component of taxes to be the residual $\tau_k^{WC} = \tau_k^* - \tau_k^M$. In Figure 6, we plot the Mirrleesian tax rate τ_k^M and the overall optimal marginal rate τ_k^* at each income percentile k and for each decade. Figure 6 shows that technical change deforms the Mirrleesian tax rate

⁴⁰GPOPS-II is a flexible software for solving optimal control problems, see Patterson and Rao (2014).

⁴¹That is, set the wage compression term to zero in (4) and rearrange. For convenience, we continue to state tax formulas and their components in their discrete, rather than continuous forms.

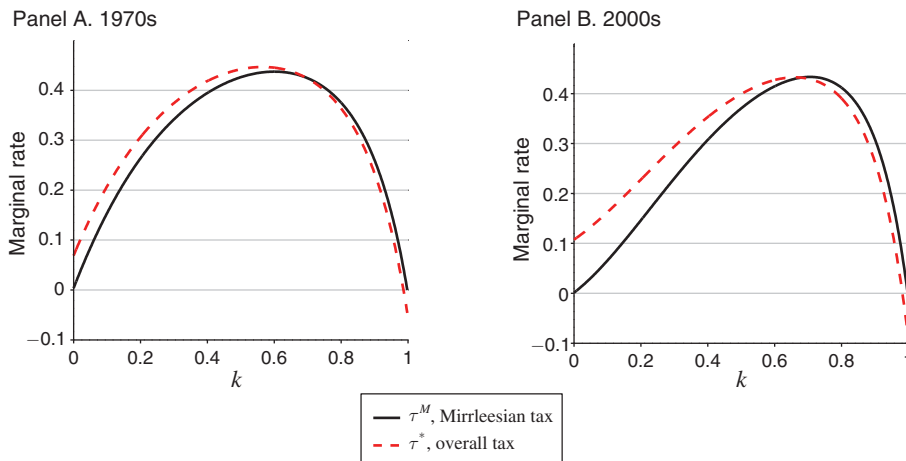


FIGURE 6. DECOMPOSING TAXES

pushing it to the right except at the lowest and highest talent. In addition, it raises the wage compression component at lower incomes and reduces it at higher ones. Overall the wage compression component becomes quantitatively more important.

Evolution of the Mirrlees Term.—We further decompose the Mirrlees term into its redistributive Ψ and wage growth parts in Figure 7.⁴² The main impact of technical change is upon wage growth (with some slight reinforcement from the redistributive term Ψ^*). This is largely driven by shifts to the a function. As noted previously, our estimates suggest that the productivities of low talents catch up with high in less complicated tasks and fall behind in more complex ones. At any effort profile and, in particular, at the optimal one, this shift compresses wage differentials at the bottom and expands them at the top. Shifts in the b function and in task demand from the middle to the extremes slightly reinforce the effect. The impact of the latter is, however, surprisingly small. This is largely because, in relevant areas of the task space, modest adjustments in the tasks of workers \tilde{v}^* are consistent with quite large variations in the density of workers across tasks \tilde{k}_v^* . Consequently, increases in the demand for low and high tasks are met with increases in the number of workers performing these tasks, but relatively little adjustment in task assignment and, hence, relative productivities and wages. For more details see online Appendix IV. The overall effect of these a and b changes is to relax incentive constraints and reduce marginal taxes at the bottom, but to tighten them and raise marginal taxes at the top.

Evolution of the Wage Compression Term.—Adjustment of the wage compression terms is in the opposite direction to the adjustment of the Mirrlees term previously described. Figure 8 displays this adjustment.

It shows that the wage compression term *rises* at low incomes and *falls*, becoming more negative, at higher ones. These changes are largely attributable to adjustments in the relative wage elasticities $\phi_{k,j}^*$. The k th talent's wage compression term is given

⁴²The other components are constant over time under our assumptions.




FIGURE 7. MIRRLEES TERM DECOMPOSITION




FIGURE 8. EVOLUTION OF THE WAGE COMPRESSION TERM

by $\Phi_k^* = \sum_{j=1}^{K-1} \mathcal{M}_{k,j} \phi_{k,j}^*$. This equation expresses Φ_k^* as a weighted sum of relative wage elasticities, with the weights depending upon the marginal incentive benefit of adjusting each pair of relative wages. Mechanically, $\phi_{k,j}^*$ is positive if $j \geq k$ and negative otherwise, so that all $\phi_{k,j}^*$ are positive if $k = 1$ and all are negative if $k = K$. For some intermediate k , positive and negative terms cancel and the wage compression term is zero. An increase in the lowest talent's effort pushes all higher

talents upwards through the task spectrum, raising the relative wages of all adjacent talents. This tightens all incentive constraints and is undesirable. Consequently, the lowest talent has the highest wage compression term and that talent's effort should be deterred at the margin. For the highest talent, this argument is reversed. An increase in the highest talent's effort pushes all lower talents downwards through the task set, compressing relative wages. This relaxes incentive constraints and should be encouraged at the margin with lower marginal taxes on high-talent incomes. For intermediate talents these effects wholly or partially offset, leading to wage compression terms that are smaller in absolute value. Figure 9 shows the impact of technical change on relative wage elasticities (normalized by population shares) $\phi_{k,j}^*/\pi_k$, $j = 1, \dots, K$ for low, mid, and high talents (labeled L , M , and H).⁴³ It indicates that almost all $\phi_{k,j}^*$ rise in absolute value. This is largely a consequence of the rise in the comparative adjustment parameter α_2 which, although it dampens the assignment response to adjustments in effort, raises the sensitivity of relative wages to any reassignment that occurs. Changes in the b function have only moderate effects on these elasticities, see online Appendix VI.




FIGURE 9. EFFORT ELASTICITY OF TASKS

Combining Terms.—The Mirrlees and wage compression terms evolve in opposite directions. Of the two, it is the adjustment to the Mirrlees term that is largest over most incomes. Consequently, marginal tax rates fall at low (but not the lowest) and rise at high (but not the highest) incomes. These adjustments are significantly muted by changes to the wage compression term and at the extremes of the wage distribution changes in this term predominate.

⁴³ Note the global impact of relative wages to an effort adjustment, an example of the ripple effect described previously.

VI. Extension: Intra-task Wage Dispersion

This paper's online appendices contain various robustness checks, extensions, and optimal tax calculations under alternative parameterizations. In the remainder of this section we focus on a specific extension that can accommodate intra-occupational wage dispersion. Recall that in the equilibrium of our (benchmark) model, differently talented workers partition the task space and all workers within a task receive the same wage: there is no intra-task wage dispersion. Our empirical strategy identifies tasks with occupations and uses dispersion in occupational average wages to determine the a function. It makes no use of measured intra-occupational wage dispersion.

Simple regressions suggest that between one-third and one-half of wage dispersion can be attributed to occupation. Mouw and Kalleberg (2010) impute wages using income and hours data in the CPS and regress this on three digit occupation dummies. They obtain an R^2 of 39 percent in the 1980s rising to 43 percent in 2010. Lane, Salmon, and Spletzer (2007) using OES microdata from 1996–1997 find that one digit occupational dummies account for 28 percent of wage variation rising to 54 percent when three digit occupational data is used. Overall, although occupations account for an important part of wage variation, significant residual wage variation remains. However, the identification of this residual variation with intra-task wage dispersion must be qualified in two ways. First, the residual absorbs measurement error in incomes and hours (from which wages are imputed).⁴⁴ Second, it absorbs occupational misclassifications and, more generally, unmeasured variation in task complexity. Several occupational categories within the CPS have fairly expansive definitions (e.g., some managerial occupations include managers of small, simple organizations, as well as managers of large complex ones) and it is likely that different workers sharing such occupational classifications perform different complexity-ranked activities.⁴⁵ It is notable that when Lane, Salmon, and Spletzer (2007) introduce establishment dummies on top of occupational ones and interact these dummies the R^2 in their regressions rises to 88 percent. While establishment dummies may capture many things, it is plausible that they help further refine the task performed by a worker (especially when interacted with occupation). To address this issue requires further unbundling of measured occupations.⁴⁶

Notwithstanding the preceding concerns, intra-task wage variation is present and does contribute to measured intra-occupational wage dispersion. We consider the extent to which it qualifies our results in online Appendix VI. We do so by extending the model and numerically parameterizing it to enhance intra-task wage dispersion. Our goal is to provide a lower bound for the responsiveness of policy to technical change. In the extended model, there are two aspects of talent: one captures comparative advantage in complex tasks, the other the ability to do all things well. Similar to our baseline model in the main text, comparative advantage types partition the

⁴⁴Bound and Krueger (1991) find that measurement error accounts for 27.6 percent of total variance of CPS earnings, while Bound et al. (1994) find that it is more severe for hours and wages.

⁴⁵Relatedly, “superstar” workers belong to (measured) occupations inhabited by much lower paid workers. It is arguable that these different workers trade in distinct task-markets with distinct shadow prices.

⁴⁶We use the O*NET to provide some very preliminary results in this direction in online Appendix VII. There we report summaries of survey results that indicate disagreement as to the knowledge requirements of occupations (amongst workers employed in or firms employing workers in these occupations). These disagreements are greatest in occupations paying higher wages.

ordered space of tasks amongst themselves. Wage variation within these partitions (and, hence, within tasks) is created by dispersion in the second (absolute advantage) component of talent. Such dispersion weakens the link between wages and tasks. It diffuses the impact of technical change and of taxes targeted at a particular income across the wage distribution. Thus, it dampens the responsiveness of policy to technical change. In taking the model to the data, we assume a coarse set of comparative advantage types and attribute all measured residual wage variation (about 75 percent of the total in our CPS sample) to variations in absolute advantage. Since the set of comparative advantage types is coarse, the partitions of the occupation space are large. Hence, we (deliberately) attribute some measured inter-occupational wage variation to absolute advantage (and as discussed above measured inter-occupational wage variation may understate inter-task wage variation and the contribution to overall wage dispersion created by the interaction of talents and tasks). As expected, the impact of technical change on marginal taxes is smaller than in our baseline case: the largest adjustment is about 2.5 points as compared to about 8.5 points before. Again, this adjustment is the net effect of countervailing changes to the Mirrlees and wage compression terms. We interpret this number as a lower bound on the responsiveness of policy to technical change. Moreover, while the quantitative response is more muted than in the benchmark case, the broad policy prescription of modest marginal tax reductions over a band of low- to mid-level incomes combined with an increase over higher incomes is robust.

VII. Conclusion

We relate the positive literature on technical change to normative work on optimal taxation by embedding an assignment model into an optimal tax framework. The assignment component induces an indirect production function over worker efforts enabling us to map technical parameters determining the productivity of task-talent matches and the demand for tasks to the variables and elasticities relevant for optimal tax analysis. We investigate the implications of changes in these parameters for optimal taxes, measure the extent of this change in US data and evaluate its implications for optimal policy.

The impacts of technical change on wage growth across talents and the substitutability of talents across tasks emerge as key drivers of policy. The twisting of the task-talent productivity function with low talents catching up in simple tasks and falling behind in more complex ones compresses wage differentials at the bottom, while expanding them at the top. It is a force for reduced marginal taxes on low incomes and increased marginal taxes on high incomes. On the other hand, increased complementarity between talent and task complexity reduces the substitutability of talents. In particular, the highest talents become increasingly locked into the highest tasks. Migration to lower ranked tasks to avoid lower task shadow prices entails greater erosion of productivity. This gives the government more tax leverage over the wage distribution. It is a force for higher marginal tax rates at the bottom. A key message of this paper is that policy depends upon the balance of these forces. Models that treat wages (or even the elasticity of substitution between talents) as exogenous omit the latter. We find its impact to be moderate, but nonnegligible.

Our paper takes a first step in integrating a task-based model of technical change into a normative public economics framework. We conclude by describing four extensions that we leave for future research. First, our model focuses on the intensive margin of labor supply.⁴⁷ It abstracts from indivisibilities in labor supply. If working at a given task requires a minimal (task specific) effort, then some workers may choose inactivity under the optimal tax code. As Saez (2002) shows such modeling of the worker extensive margin can significantly affect optimal tax results at the bottom. However, its implications for the impact of technical change on tax design are less clear. Second, our model assumes that the matching of talents to tasks is frictionless. Thus, our quantitative work is best viewed as capturing the long run policy response to technical change after the (possibly slow) reassignment of workers to tasks following such change. The role of income taxation in supplementing other sources of insurance during transitions is omitted. Third, our model omits accumulation of experience or skill within tasks that can impede or promote transitions to other tasks. Fourth, we abstract from the endogenous nature of technical change. Relaxing these restrictions remain important topics for further research.

APPENDIX

It is straightforward to verify that any allocation that solves the mechanism design problem (MDP) is implementable as part of a tax equilibrium. On the other hand, the allocation from a tax equilibrium is feasible for (MDP). Consequently, an optimal tax equilibrium may be constructed from a solution to (MDP) $\{c_k^*, e_k^*\}_{k=1}^K$ by associating with it the wages and taxes needed to implement this solution.

We make two preliminary observations on solutions to (MDP). First, given a solution $\{c_k^*, e_k^*\}_{k=1}^K$, worker types may be ordered according to their optimal shadow wages $\{w_k^*\}_{k=1}^K$. Types whose wages are tied may be further ordered by their pre-tax incomes $q_k^* = w_k^* e_k^*$. Types may then be relabeled accordingly (with ties between both wage and income ordered arbitrarily). Thus, the k th worker type has a wage that is weakly greater than the wages of types 1 to $k - 1$ and if the k th type's wage ties with the $k - 1$ th type, then its income is weakly greater. We impose this labeling below. Second, only a subset of incentive constraints (2) bind. Recall that a (k, j) th incentive constraint is local if $j \in \{k - 1, k + 1\} \cap \{1, \dots, K\}$; otherwise it is non-local. A well known consequence of the Spence-Mirrlees single crossing property and the structure of the incentive constraints in settings with exogenous wages is that non-local incentive constraints do not bind at an optimum. This result continues to hold in the present setting under our ordering.⁴⁸ We record this fact in Lemma 2.

⁴⁷ However, our model admits an alternative interpretation in which workers exert effort in skill accumulation rather than market work. Our theoretical insights are applicable to this interpretation.

⁴⁸ We omit the proof. It follows from a slight modification of Theorems 3 and 4 in Milgrom and Shannon (1994). In our setting it is possible for two worker types k and $k + 1$ to have the same wage, but different efforts, incomes, and consumptions at the optimum. If the k th type has a higher income than the $k + 1$ th type, then it is possible that the (nonlocal) $(k - 1, k + 1)$ and $(k, k + 2)$ incentive constraints bind. Thus, ordering of worker types with tied wages by income is necessary to ensure only local incentive constraints bind.

LEMMA 2: Let $\{c_k^*, e_k^*\}_{k=1}^K$ denote a solution to (MDP) with corresponding shadow wages $\{w_k^*\}_{k=1}^K$, $w_k^* = F_k(e_1^* \pi_1, \dots, e_K^* \pi_K)$ and incomes $q_k^* = w_k^* e_k^*$ and types labeled consistently with their ranking in the wage (and when wages are tied income) distribution, then (i) $c_{k+1}^* \geq c_k^*$ and $q_{k+1}^* \geq q_k^*$, and (ii) nonlocal incentive constraints do not bind.

As in the main text, we gather the constraint functions from the incentive and resource constraints into the single function $\mathcal{G}: \mathbb{R}_+^{2K} \rightarrow \mathbb{R}^{K(K-1)}$ and say that $\{c_k, e_k\}_{k=1}^K \in \mathbb{R}_{++}^{2K}$ satisfies the constraint qualification if there is an $x \in \mathbb{R}^{2K}$ such that $\nabla \mathcal{G}(\{c_k, e_k\}_{k=1}^K)x < 0$.

PROPOSITION 3: Let T^* and $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$ denote an optimal tax equilibrium with worker types indexed so that $w_k^* = F_k(e_1^* \pi_1, \dots, e_K^* \pi_K)$ is non-decreasing in k . Assume that $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$ is interior (i.e., in \mathbb{R}_{++}^{2K}) and that \mathcal{G} satisfies the constraint qualification at $\{c_k^*, e_k^*, w_k^*\}_{k=1}^K$, then optimal tax rates satisfy:

$$(27) \quad \frac{\tau_k^*}{1 - \tau_k^*} = \underbrace{\frac{1 - \Pi_k}{\pi_k} \left\{ \frac{\Delta w_{k+1}^*}{w_{k+1}^*} \Psi_{k,k+1}^* \mathcal{H}_{k,k+1}^* - \frac{\Delta w_k^*}{w_{k-1}^*} \Psi_{k,k-1}^* \mathcal{H}_{k,k-1}^* \right\}}_{\text{Mirrlees}} + \underbrace{\sum_{j=1}^{K-1} \mathcal{M}_{k,j}^* \phi_{k,j}^*}_{\text{Wage compression}}$$

where $\Delta w_k^* := w_k^* - w_{k-1}^*$, $\Psi_{k,k+1}^* := \frac{U_c(c_k^*, e_k^*)}{1 - \Pi_k} \frac{\eta_{k+1,k}^*}{\chi^*}$ and $\Psi_{k,k-1}^* := \frac{U_c(c_k^*, e_k^*)}{1 - \Pi_k} \frac{\eta_{k-1,k}^*}{\chi^*}$ are normalized optimal multipliers on the $(k+1, k)$ th and $(k-1, k)$ th incentive constraints,

$$\begin{aligned} \mathcal{H}_{k,j}^* &:= - \frac{\Delta_e U_c(c_k^*, e_k^*, q_k^*/w_j^* - e_k^*)}{U_c(c_k^*, e_k^*)} e_k^* \\ &+ \frac{\Delta_e U_e(c_k^*, e_k^*, q_k^*/w_j^* - e_k^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_j^*} e_k^* + 1, \end{aligned}$$

with $q_k^* = w_k^* e_k^*$,

$$\mathcal{M}_{k,j}^* :=$$

$$\frac{U_c(c_k^*, e_k^*)}{U_e(c_k^*, e_k^*) e_k^*} \left[\frac{\eta_{j+1,j}^*}{\chi^*} U_e \left(c_j^*, \frac{q_j^*}{w_{j+1}^*} \right) \frac{q_j^*}{w_{j+1}^*} - \frac{\eta_{j,j+1}^*}{\chi^*} U_e \left(c_{j+1}^*, \frac{q_{j+1}^*}{w_j^*} \right) \frac{q_{j+1}^*}{w_j^*} \right] \frac{1}{\pi_k},$$

$$\text{and } \phi_{k,j}^* = \frac{e_k^*}{w_{j+1}^*/w_j^*} \frac{\partial w_{j+1}^*/w_j^*}{\partial e_k} (e_1^*, \dots, e_K^*).$$

PROOF:

By the preceding discussion if $\{T^*, \{c_k^*, e_k^*, w_k^*\}_{k=1}^K\}$ is an optimal tax equilibrium, then $\{c_k^*, e_k^*\}_{k=1}^K$ solves (MDP). Since \mathcal{G} satisfies the constraint qualification at $\{c_k^*, e_k^*\}_{k=1}^K$ and $\{c_k^*, e_k^*\}_{k=1}^K$ is interior to \mathbb{R}_{++}^{2K} , then $\{c_k^*, e_k^*\}_{k=1}^K$ satisfies Karush-Kuhn-Tucker conditions with multipliers χ^* and $\eta_{k,j}^*$ on the resource and incentive constraints (and zero multipliers on the nonnegativity conditions $c_k^*, e_k^* \geq 0$). Also, since worker types are indexed so that shadow wages $w_k^* = F_k(e_1^* \pi_1, \dots, e_K^* \pi_K)$ are non-decreasing in k and if wages of different types are tied so that incomes are non-decreasing in k , then by Lemma 2, only local incentive constraints are potentially binding and, hence, only the $\eta_{k,k-1}^*$ and $\eta_{k,k+1}^*$ multipliers are potentially nonzero. The first order condition for e_k^* reduces to

$$-U_e(c_k^*, e_k^*) = \chi^* w_k^* \pi_k / D_k^*,$$

where:

$$D_k^* := g_k + \eta_{k,k-1}^* - \eta_{k+1,k}^* \frac{U_e(c_k^*, q_k^* / w_{k+1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_{k+1}^*} + \eta_{k,k+1}^* - \eta_{k-1,k}^* \frac{U_e(c_k^*, q_k^* / w_{k-1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_{k-1}^*} \\ + \frac{\chi^* \pi_k}{U_c(c_k^*, e_k^*)} \Phi_k^* + \frac{\chi^* \pi_k}{U_c(c_k^*, e_k^*)} \Upsilon_k^* \text{ and}$$

$$\Phi_k^* := \frac{U_c(c_k^*, e_k^*)}{\pi_k} \sum_{j=1}^{K-1} \frac{\eta_{j+1,j}^*}{\chi^*} \frac{U_e(c_j^*, q_j^* / w_{j+1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_j^* e_j^*}{w_{j+1}^* e_k^*} \phi_{k,j}^*,$$

and

$$\Upsilon_k^* := \frac{U_c(c_k^*, e_k^*)}{\pi_k} \sum_{j=1}^{K-1} \frac{\eta_{j-1,j}^*}{\chi^*} \frac{U_e(c_j^*, q_j^* / w_{j-1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_j^* e_j^*}{w_{j+1}^* e_k^*} \phi_{k,j-1}^*.$$

The first order condition for c_k^* reduces to

$$U_c(c_k^*, e_k^*) = \frac{\chi^* \pi_k}{g_k + \eta_{k,k-1}^* - \eta_{k+1,k}^* \frac{U_e(c_k^*, q_k^* / w_{k+1}^*)}{U_e(c_k^*, e_k^*)} + \eta_{k,k+1}^* - \eta_{k-1,k}^* \frac{U_e(c_k^*, q_k^* / w_{k-1}^*)}{U_e(c_k^*, e_k^*)}}.$$

Define the consumption-effort wedge: $\frac{\tau_k^*}{1 - \tau_k^*} = -\frac{w_k^* U_c(c_k^*, e_k^*)}{U_e(c_k^*, e_k^*)} - 1$. Combining expressions gives

$$\frac{\tau_k^*}{1 - \tau_k^*} = \frac{U_c(c_k^*, e_k^*)}{\pi_k} \left\{ \frac{\eta_{k+1,k}^*}{\chi^*} \left\{ \frac{U_e(c_k^*, q_k^* / w_{k+1}^*)}{U_e(c_k^*, e_k^*)} - \frac{U_e(c_k^*, q_k^* / w_{k+1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_{k+1}^*} \right\} \right. \\ \left. + \frac{\eta_{k-1,k}^*}{\chi^*} \left\{ \frac{U_e(c_k^*, q_k^* / w_{k-1}^*)}{U_e(c_k^*, e_k^*)} - \frac{U_e(c_k^*, q_k^* / w_{k-1}^*)}{U_e(c_k^*, e_k^*)} \frac{w_k^*}{w_{k-1}^*} \right\} \right\} + \Phi_k^* + Y_k^*.$$

The formulas in Proposition 3 then follow immediately from the definitions of Φ_k^* , Y_k^* , $\mathcal{H}_{k,j}^*$, $\mathcal{M}_{k,j}^*$, and $\phi_{k,j}^*$ after substitution into and rearrangement of the preceding expression. ■

The terms on the right-hand-side of the optimal tax formula in Proposition 3 are generalizations of the “Mirrlees” and “Wage Compression” terms obtained in the main text. These terms incorporate the impact of binding (local) upwards constraints as well as downwards constraints. In standard models upwards constraints bind when it is optimal to pool agents with *distinct wages* at a *common consumption-effort allocation*. In the more general problem, they may also bind when it is optimal to pool distinct types with *distinct allocations* at a *common wage*. Rothschild and Chen (2014) provide an example in which such wage pooling occurs.⁴⁹ Our later assignment model micro-founds the production function F . In that setting, the induced production function does not feature wage pooling. This motivates us to consider situations in which the local upwards incentive constraints $(k, k + 1)$ are strictly non-binding at the optimum:⁵⁰

$$(\text{NUIC}) \quad U(c_k^*, e_k^*) > U(c_{k+1}^*, q_{k+1}^*/w_k^*).$$

In such cases, the optimal tax formula (27) reduces to that given in Proposition 1. The latter is obtained as a simple corollary of Proposition 3.

PROOF OF PROPOSITION 1:

The optimal tax formula (4) in Proposition 1 follows directly from that in Proposition 3 after setting all $\eta_{k-1,k}^*$ equal to 0, using the modified definitions in Proposition 1 and expanding the recursion for $\eta_{k+1,k}^*$ implied by the first order condition for c_{k+1}^* :

$$\frac{\eta_{k+1,k}^*}{\chi^*} = \left(1 - \frac{g_{k+1} U_c(c_{k+1}^*, e_{k+1}^*)}{\chi^* \pi_{k+1}}\right) \frac{\pi_{k+1}}{U_c(c_{k+1}^*, e_{k+1}^*)} + \frac{\eta_{k+2,k+1}^*}{\chi^*} \frac{U_c(c_{k+1}^*, q_{k+1}^*/w_{k+2}^*)}{U_c(c_{k+1}^*, e_{k+1}^*)},$$

with $\eta_{K+1,K}^* = 0$. ■

REFERENCES

- Acemoglu, Daron.** 2002. “Directed Technical Change.” *Review of Economic Studies* 69 (4): 781–809.
- Acemoglu, Daron, and David Autor.** 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*, Vol. 4B, edited by David Card and Orley Ashenfelter, 1043–1171. Amsterdam: Elsevier B. V.
- Ales, Laurence, Musab Kurnaz, and Christopher Sleet.** 2015. “Technical Change, Wage Inequality, and Taxes: Dataset.” *American Economic Review*. <http://dx.doi.org/10.1257/aer.20140466>.
- Autor, David H., and David Dorn.** 2013. “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.” *American Economic Review* 103 (5): 1553–97.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney.** 2006. “The Polarization of the US Labor Market.” *American Economic Review* 96 (2): 189–94.
- Autor, David H., Lawrence F. Katz, and Alan B. Krueger.** 1998. “Computing Inequality: Have Computers Changed the Labor Market?” *Quarterly Journal of Economics* 113 (4): 1169–1213.
- Autor, David H., Frank Levy, and Richard J. Murnane.** 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *Quarterly Journal of Economics* 118 (4): 1279–1333.

⁴⁹The effort of two worker types k and k' sharing a common wage may interact differently with that of a third worker type \hat{k} . Thus, it may be desirable to give k and k' distinct allocations (lying upon the same indifference curve).

⁵⁰Standard pooling of consumption-effort allocations across types with distinct wages is still, in principle, possible. Inequality (NUIC) excludes this.

- Bound, John.** 1994. "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics* 12 (3): 345–68.
- Bound, John, and Alan B. Krueger.** 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics* 9 (1): 1–24.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt.** 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117 (1): 339–76.
- Buera, Francisco J., and Joseph P. Kaboski.** 2012. "The Rise of the Service Economy." *American Economic Review* 102 (6): 2540–69.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber.** 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3): 471–75.
- Costinot, Arnaud, and Jonathan Vogel.** 2010. "Matching and Inequality in the World Economy." *Journal of Political Economy* 118 (4): 747–86.
- Diamond, Peter, Daniel McFadden, and Miguel Rodriguez.** 1978. "Measurement of the Elasticity of Factor Substitution and Bias of Technical Change." In *Production Economics: A Dual Approach to Theory and Applications*, edited by Melvyn Fuss and Daniel McFadden, 125–47. Amsterdam: Elsevier.
- Goldin, Claudia, and Lawrence F. Katz.** 1998. "The Origins of Technology-Skill Complementarity." *Quarterly Journal of Economics* 113 (3): 693–732.
- Goos, Maarten, and Alan Manning.** 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics* 89 (1): 118–33.
- Grossman, Gene M., and Esteban Rossi-Hansberg.** 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review* 98 (5): 1978–97.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2014. "Optimal Tax Progressivity: An Analytical Framework." Federal Reserve Bank of Minneapolis Staff Report 496.
- Hellwig, Martin F.** 2007. "The Undesirability of Randomized Income Taxation under Decreasing Risk Aversion." *Journal of Public Economics* 91 (3–4): 791–816.
- Kaplan, Steven N., and Joshua Rauh.** 2013. "It's the Market: The Broad-Based Rise in the Return to Top Talent." *Journal of Economic Perspectives* 27 (3): 35–56.
- King, Miriam, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick.** 2010. *Integrated Public Use Microdata Series, Current Population Survey: Version 3.0*. [Machine-readable database]. Minneapolis: University of Minnesota.
- Krusell, Per, Lee E. Ohanian, José-Victor Ríos-Rull, and Giovanni L. Violante.** 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–53.
- Lane, Julia I., Laurie A. Salmon, and James R. Spletzer.** 2007. "Establishment Wage Differentials." *Monthly Labor Review* 130 (4): 3–17.
- Lockwood, Benjamin B., Charles G. Nathanson, and E. Glen Weyl.** 2014. "Taxation and the Allocation of Talent." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1324424 (accessed January 26, 2015).
- Milgrom, Paul, and Chris Shannon.** 1994. "Monotone Comparative Statics." *Econometrica* 62 (1): 157–80.
- Mirrlees, James A.** 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (114): 175–208.
- Mouw, Ted, and Arne L. Kalleberg.** 2010. "Occupations and the Structure of Wage Inequality in the United States, 1980s to 2000s." *American Sociological Review* 75 (3): 402–31.
- Patterson, Michael A., and Anil V. Rao.** 2014. "GPOPS-II: A MATLAB Software for Solving Multiple-Phase Optimal Control Problems Using hp-Adaptive Gaussian Quadrature Collocation Methods and Sparse Nonlinear Programming." *ACM Transactions on Mathematical Software* 41 (1). <http://dx.doi.org/10.1145/2558904>.
- Rothschild, Casey, and Lizi Chen.** 2014. "Screening with Endogenous Preferences." Unpublished.
- Rothschild, Casey, and Florian Scheuer.** 2013. "Redistributive Taxation in the Roy Model." *Quarterly Journal of Economics* 128 (2): 623–68.
- Rothschild, Casey, and Florian Scheuer.** 2014. "A Theory of Income Taxation under Multidimensional Skill Heterogeneity." National Bureau of Economic Research Working Paper 19822.
- Roy, A. D.** 1950. "The Distribution of Earnings and of Individual Output." *Economic Journal* 60 (239): 489–505.
- Saez, Emmanuel.** 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies* 68 (1): 205–29.
- Saez, Emmanuel.** 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics* 117 (3): 1039–73.

- Salanié, Bernard.** 2011. *Economics of Taxation*. Cambridge, MA: MIT Press.
- Sattinger, Michael.** 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica* 43 (3): 455–68.
- Slavík, Ctirad, and Hakki Yazici.** 2014. "Machines, Buildings, and Optimal Dynamic Taxes." *Journal of Monetary Economics* 66: 47–61.
- Stiglitz, Joseph E.** 1982. "Self-Selection and Pareto Efficient Taxation." *Journal of Public Economics* 17 (2): 213–40.
- Teulings, Coen N.** 1995. "The Wage Distribution in a Model of the Assignment of Skills to Jobs." *Journal of Political Economy* 103 (2): 280–315.