Robots, Trade, and Luddism *

Arnaud Costinot
MIT

Iván Werning
MIT

June 2018

Abstract

Technological change, from the advent of robots to expanded trade opportunities, tends to create winners and losers. When are such changes welcome? How should government policy respond? We consider these questions in a second best world, with a restricted set of tax instruments. Our first set of results shows that, despite these restrictions, productivity improvements are always welcome and valued in the same way as in a first best world. Our second set of results offers optimal tax formulas, as well as bounds on those taxes, that can be computed using existing reduced-form evidence on the distributional impact of robots and trade. Our final result shows that while distributional concerns create a rationale for non-zero taxes on robots and trade, more robots, more trade, and more inequality may be optimally met with lower taxes.

*We thank Andres Rodriguez-Clare for detailed comments on an earlier draft, seminar and conference participants at UC Berkeley, Harvard and BFI for helpful feedback, Masao Fukui, Mary Gong and Martina Uccioli for excellent research assistance, and HAL for its unrelenting discouragement.

1 Introduction

Robots and artificial intelligence technologies are on the rise. So are imports from China and other developing countries. Regardless of its origins, technological change creates opportunities for some workers, destroys opportunities for others, and generates significant distributional consequences, as documented in the recent empirical work of Autor, Dorn and Hanson (2013) and Acemoglu and Restrepo (2017) for the United States.

When should technological change be welcome? Should any policy response be in place? And if so, how should we manage new technologies? Should we become more luddites as machines become more efficient or more protectionist as trade opportunities expand? The goal of this paper is to provide a general second-best framework to help address these and other related questions.

In seeking answers to these questions one first needs to take a stand on the range of available policy instruments. Obviously, if the idilic lump-sum transfers are available, distribution can be done efficiently, without distorting production. Even in the absence of lump-sum transfers, if linear taxes are available on all goods and factors, production efficiency may hold, as in Diamond and Mirrlees (1971b). In both cases, zero taxes on robots and free trade are optimal. At another extreme, in the absence of any policy instrument, whenever technological progress creates at least one loser, a welfare criterion must be consulted and the status quo may be preferred.

Here, we focus on intermediate, and arguably more realistic, scenarios where tax instruments are available, but are more limited than those ensuring production efficiency. Our framework is designed to capture general forms of technological change. We consider two sets of technologies, which we refer to as "old" and "new". For instance, firms using the new technology may be producers of robots or traders that export some goods in exchange for others. Since we are interested in the optimal regulation of the new technology, we do not impose any restriction on the taxation of firms using that technology, e.g. taxes on robots or trade. In contrast, to allow for a meaningful trade-off between redistribution and efficiency, we restrict the set of taxes that can be imposed on firms using the old technology as well as on consumers and workers. In the economic environment that we consider, the after-tax wage structure can be influenced by tax policy, but not completely controlled.

Our first set of results focuses on the welfare impact of new technologies under the assumption that constrained, but optimal policies are in place. We offer a novel envelope result that generalizes the evaluation of productivity shocks in first-best environments, as in Solow (1957) and Hulten (1978), to distorted economies. Because of restrictions on the set of available tax instruments, marginal rates of substitution may not be equalized across agents and marginal rates of transformation may not be equalized between new and old technology firms. Yet, "Immiserizing Growth," as in Bhagwati (1958), never arises. Provided that governments can tax new technology firms, the welfare impact of technological progress can be measured in the exact same way as in first best environments (despite not being first best).

A direct implication of our envelope result is that even if new technologies tend to have a disproportionate effect on the wages of less skilled workers, and we care about redistribution, this does not create any new rationale for taxes and subsidies on innovation. In the case of terms-of-trade shocks, our envelope result states that such shocks are beneficial if and only if they raise the value of the trade balance at current quantities. In other words, the gains from international trade can still be computed by integrating below the demand curves for foreign goods. Distributional considerations may affect how much we trade, but not the mapping between observed trade flows and welfare, as in Arkolakis, Costinot and Rodríguez-Clare (2012).

Our second set of results characterizes the structure of Pareto efficient taxes on old and new technology firms in general environments when the taxation of different factors of production is limited to non-linear income taxation, as in Mirrlees (1971). The work of Naito (1999) has proven that governments seeking to redistribute income from high- to low-skill workers may have incentives to depart from production efficiency. Doing so manipulates relative wages, which cannot be taxed directly, and relaxes incentive compatibility constraints. Our analysis generalizes this result and and goes beyond qualitative insights by deriving optimal tax formulas as well as bounds on any Pareto efficient tax.

We offer our sharpest results in the context of an economy with one final good and one intermediate good. The intermediate good could be either robots, that are produced domestically, or other machines, that are imported from abroad in exchange for the final good. We set up this environment with limited general equilibrium effects. Namely, we restrict technology so that the number of machines may directly affect inequality by affecting relative marginal products of labor, but not indirectly through further changes in relative labor supply, as in Stiglitz (1982).

We show that the Pareto efficient tax on machines can be expressed in terms of a few sufficient statistics:factor shares, marginal income tax rates, labor supply elasticities, and the elasticity of relative wages with respect to the number of machines. Crucially, our formula does not require knowledge, or assumptions, about the Pareto weights assigned to different agents in societies. Such considerations are implicitly revealed by the observed income tax schedule. Using the reduced-form evidence of Acemoglu and Restrepo (2017) on the impact of robots and the evidence of Chetverikov, Larsen and Palmer (2016) on the impact of Chinese imports on inequality as illustrations, we then demonstrate how to compute the Pareto efficient tax on robots and Chinese imports, respectively.

We conclude our analysis with a comparative static exercise that asks: as progress in Artificial Intelligence makes for cheaper and better robots, should we tax them more? Or in a trade context, does hyper-globalization call for hyper-protection? Through a simple parametric example, we show that in contrast to many popular discussions, improvements in new technologies may be associated with more robots and more trade, as well as more inequality caused by those improvements, but lower Pareto efficient taxes on robots and trade.

Our paper makes three distinct contributions to the existing literature. The first one is a new perspective on the welfare impact of technological progress in the presence of distortions. In a first best world, the impact of small productivity shocks can be evaluated, absent any restriction on preferences and technology, using a simple envelope argument as in Solow (1957) and Hulten (1978). With distortions, evaluating the welfare impact of productivity shocks, in general, requires additional information about whether such shocks aggravate or alleviate underlying distortions. In an environment with markups, for instance, this boils down to whether employment is reallocated towards goods with higher or lower markups, as in Basu and Fernald (2002), Arkolakis, Costinot, Donaldson and Rodríguez-Clare (Forthcoming), and Baeqee and Farhi (2017). If the aggravation of distortions is large enough, technological progress may even lower welfare, as discussed by Bhagwati (1971). Here, we follow a different approach. Our analysis builds on the idea that while economies may be distorted and tax instruments may be limited, the government may still have access to policy instruments to control the new technology. If so, the envelope results of Solow (1957) and Hulten (1978) still hold, with direct implications for the measurement of the welfare gains from globalization and automation as well as for the taxation of innovation.

Our second contribution is a general characterization of the structure of Pareto efficient taxes in environments with restricted factor income taxation. In so doing, we fill a gap between the general analysis of Diamond and Mirrlees (1971b,a) and Dixit and Norman (1980), which assumes that linear taxes on all factors are available, and specific examples, typically with two goods and two factors, in which only income taxation is available, as in the original work of Naito (1999), and subsequent work by Guesnerie (1998), Spector (2001), and Naito (2006).[1] On the broad spectrum of restrictions on avail-

[1]In all three papers, like in Dixit and Norman (1980), the new technology is international trade. In another related trade application, Feenstra and Lewis (1994) study an environment where governments

able policy instruments, one can also view our analysis as an intermediate step between the work of Diamond and Mirrlees (1971b,a) and Dixit and Norman (1980) and the trade policy literature, as reviewed for instance in Rodrik (1995), where it is common to assume that the only instruments available for redistribution are trade taxes.

Our third contribution is to offer a more specific application of our general formulas to the taxation of robots and trade. In recent work, Guerreiro, Rebelo and Teles (2017) have studied a model with both skilled and unskilled workers as well as robots, with a nested CES production function. Assuming factor-specific taxes are unavailable, they find a nonzero tax on robots to be generally optimal, in line with Naito (1999). Although we share the same rationale for finding nonzero taxes on robots, based on Naito (1999), our main goal is not to sign the tax on robots, nor to explore a particular production structure, but instead to offer tax formulas highlighting key sufficient statistics needed to determine the level of taxes, with fewer structural assumptions. In this way, our formulas provide a foundation for existing empirical work as well as the basis for novel comparative static results.[2] In another recent contribution, Hosseini and Shourideh (2018) analyze a multi-country Ricardian model of trade with input-output linkages and imperfect mobility of workers across sectors. Although sector-specific taxes on labor are not explicitly allowed, these missing taxes can be perfectly mimicked by the available tax instruments. By implication, their economy provides an alternative implementation but fits Diamond and Mirrlees (1971a,b) and Dixit and Norman (1980, 1986), where households face a complete set of linear taxes, including sector-specific taxes on labor. Production efficiency and free trade then follow, just as they did in Diamond and Mirrlees (1971a,b).[3]

2 Environment

We consider an economy with a finite number of goods indexed by $i = 1, ..., N$ and a continuum of households indexed by their ability $\theta \in [\underline{\theta}, \bar{\theta}]$. We let F denote the cumulative distribution of abilities in the population and f its density.

cannot subject different worker types to different taxes, but can offer subsidies to workers moving from one industry to another in response to trade. They provide conditions under which such a trade adjustment assistance program are sufficient to guarantee Pareto gains from trade, as in Dixit and Norman (1980).

[2]Our specific tax formulas in Section 6 relate to the work of Jacobs (2015) who considers the same tax instruments as in Naito (1999), in an environment without robots, featuring a continuum of workers and no general equilibrium effects.

[3]A separate line of work, e.g. Itskhoki (2008), Antras, de Gortari and Itskhoki (2017) and Tsyvinski and Werquin (2018), studies technological changes such as trade or robots, without considering taxes on these new technologies, but isntead focusing on how the income tax schedule may respond to these changes.

2.1 Preferences

Households have identical and weakly separable preferences between goods and labor. The utility of household θ is given by

$$U(\theta) = u(C(\theta), n(\theta)),$$
$$C(\theta) = v(c(\theta)),$$

where $C(\theta)$ is the sub-utility that household θ derives from consuming goods, $n(\theta)$ is her labor supply, $c(\theta) \equiv \{c_i(\theta)\}$ is her vector of good consumption, and u and v are the utility functions associated with lower- and upper-level preferences, respectively.

2.2 Technology

There are two types of technologies, which we refer to as "old" and "new," each associated with a distinct production set. In our applications, the new technology may represent trade with the rest of the world or the production of machines, like robots. This dichotomy between new and old technology, as opposed to the consolidation into a single aggregate production set, allows us to consider the differential taxation of firms using these two technologies, and, in turn, allows for aggregate production inefficiency.

Old Technology. Let $y \equiv \{y_i\}$ denote the vector of total net output by old technology firms and let $n \equiv \{n(\theta)\}$ denote the schedule of their total labor demand. Positive values for y_i represent output, while negative y_i represent inputs. The production set associated with the old technology corresponds to all production plans (y, n) such that

$$G(y, n) \leq 0,$$

where G is some convex and homogeneous function of (y, n). Homogeneity of G implies constant returns to scale.

New Technology. Let $y^* \equiv \{y_i^*\}$ denote the vector of total net output by new technology firms. The production set associated with the new technology corresponds to all production plans y^* such that

$$G^*(y^*; \phi) \leq 0,$$

where G^* is some convex and homogeneous function of y^* and $\phi > 0$ is a productivity parameter. We assume that G^* is decreasing in ϕ so that an increase in ϕ corresponds to an improvement in the new technology.

Unlike the old technology, the new technology abstracts from employment of labor. This is both convenient and well suited to our two main applications: trade and robots. In the first case, new technology firms are traders who can export and import goods,

$$G^*(y^*;\phi) = \bar{p}(\phi) \cdot y^*,$$

where $\bar{p}(\phi) \equiv \{\bar{p}_i(\phi)\}$ denotes the vector of world prices and $\cdot$ denotes the inner product of two vectors. An increase in ϕ corresponds to a positive terms-of-trade shock, that may be due to a change in transportation costs or productivity in the rest of the world. In the second case, new technology firms may be robot-producers that transform a composite of all other goods in the economy, call it gross output, into robots. Abstracting from labor in the new technology is also convenient, as it implies that wages are determined by the old technology. New technology has an effect on wages, through its effect on the structure of production within the old technology, but not directly through employment.

Resource Constraint. For all goods, the demand by households is equal to the supply by old and new technology firms,

$$\int c(\theta) dF(\theta) = y + y^*. \tag{1}$$

2.3 Prices and Taxes

Factors. Let $w \equiv \{w(\theta)\}$ denote the schedule of wages faced by firms. Because of income taxation, a household with ability θ and labor earnings $x(\theta) \equiv w(\theta)n(\theta)$ retains

$$R(x(\theta);\theta) = x(\theta) - T(x(\theta);\theta),$$

where $T((x(\theta);\theta)$ denotes its total tax payment. In general, the previous specification may allow for lump-sum transfers, if $T(x;\theta) = T(\theta)$, linear factor taxation, if $T(x;\theta) = (1 - T(\theta))x$, or anonymous income taxation, if $T(x;\theta) = T(\theta)$. For now, we assume that the government may be limited in its ability to tax labor income through T. This, in turn, imposes restrictions on the set of feasible retention functions,

$$\mathcal{R} \equiv \{R | R(x;\theta) = x - T(x;\theta) \text{ for all } x, \theta \text{ for some feasible } T\}.$$

In Sections 5 and 6, we will focus on the case with anonymous income taxation.

Goods. Let $p \equiv \{p_i\}$ and $p^* \equiv \{p_i^*\}$ denote the vector of good prices faced by old and new technology firms, respectively. Because of ad-valorem taxes $t \equiv \{t_i\}$ and $t^* \equiv \{t_i^*\}$, these prices may differ from the vector of good prices $q \equiv \{q_i\}$ faced by households,

$$q_i = (1 + t_i)p_i, \text{ for all } i, \tag{2}$$

$$q_i = (1 + t_i^*)p_i^*, \text{ for all } i. \tag{3}$$

Production inefficiency arises if $t \neq t^*$. In a trade context, an import tariff or an export subsidy on good i corresponds to $t_i^* > 0$, whereas an import subsidy or an export tax corresponds to $t_i^* < 0$. Likewise, a production subsidy on domestic firms corresponds to $t_i < 0$ and a tax to $t_i > 0$.

We assume that taxes on new technology firms are unrestricted, $t^* \in [-1, \infty)^N$. Hence, the government can freely choose the vector of prices faced by new technology firms, $p^* \in \mathbb{R}_+^N$, irrespectively of what the prices faced by old technology firms and households, p and q, may be. In contrast, we assume that the government may be limited in its ability to tax old technology firms through t, which imposes restrictions on the set of feasible prices that households and old technology firms may face,

$$\mathcal{P} \equiv \{(p, q) | q_i = (1 + t_i)p_i \text{ for all } i \text{ for some feasible } t\}.$$

In Sections 5 and 6, we will focus on the two extreme cases with unrestricted taxes on old technology firms, $t \in [-1, \infty)^N$, and without any taxes on those firms, $t \in \{0\}$.

3 Competitive Equilibrium with Pareto Efficient Taxes

3.1 Competitive Equilibrium

Demand. Households maximize utility taking prices and taxes as given. Since preferences are weakly separable, the demand of any household θ is given by the solution to the following two-step problem

$$c(\theta) \in \text{argmin}_{\tilde{c}(\theta)}\{q \cdot \tilde{c}(\theta) | v(\tilde{c}(\theta)) \geq C(\theta)\}, \tag{4}$$

$$C(\theta), n(\theta) \in \text{argmax}_{\tilde{C}(\theta), \tilde{n}(\theta)}\{u(\tilde{C}(\theta), \tilde{n}(\theta)) | e(q, \tilde{C}(\theta)) \leq R(w(\theta)\tilde{n}(\theta); \theta)\}, \tag{5}$$

where $e(q, C(\theta)) \equiv \min_{\tilde{c}(\theta)}\{q \cdot \tilde{c}(\theta) | v(\tilde{c}(\theta)) \geq C(\theta)\}$ denotes the expenditure function associated with the lower stage.

Supply. Firms maximize profits taking prices and taxes as given,

$$y, n \in \operatorname{argmax}_{\tilde{y},\tilde{n}}\{p \cdot \tilde{y} - \int w(\theta)\tilde{n}(\theta)dF(\theta) | G(\tilde{y},\tilde{n}) \leq 0\}, \tag{6}$$

$$y^* \in \operatorname{argmax}_{\tilde{y}^*}\{p^* \cdot \tilde{y}^* | G^*(\tilde{y}^*;\phi) \leq 0\}. \tag{7}$$

Equilibrium. A competitive equilibrium with taxes (T, t, t^*) corresponds to quantities, $c \equiv \{c(\theta)\}$, $n \equiv \{n(\theta)\}$, $C \equiv \{C(\theta)\}$, $y \equiv \{y_i\}$, and $y^* \equiv \{y_i^*\}$, and prices, $w \equiv \{w(\theta)\}$, $p \equiv \{p_i\}$, $p^* \equiv \{p_i^*\}$, and $q \equiv \{q_i\}$ such that:

i. good markets clear, condition (1);

ii. prices satisfy the non-arbitrage conditions (2) and (3);

iii. households maximize their utility, condition (4) and (5);

iv. firms maximize their profits, conditions (6) and (7).

Two comments are in order. First, since all markets clear and all budget constraints hold, the government budget constraint necessarily holds in a competitive equilibrium, an expression of Walras' law. Second, since demand and supply only depend on relative prices, we can normalize prices and taxes such that $p_1 = p_1^* = q_1 = 1$ and $t_1 = t_1^* = 0$. We will maintain this normalization throughout our paper.

3.2 Pareto Efficient Taxes

We focus on Pareto efficient taxes. That is, we assume that the government chooses a competitive equilibrium with feasible taxes (T, t, t^*) in order to maximize

$$\int U(\theta) d\Lambda(\theta),$$

where Λ denotes the distribution of Pareto weights. Λ is positive, increasing, right-continuous, and normalized so that $\Lambda(\bar{\theta}) = 1$. The utilitarian benchmark corresponds to $\Lambda = F$, the Rawlsian benchmark to $\Lambda(\theta) = 1$ for all θ, that is, full weight at $\underline{\theta}$.[4]

To facilitate the analysis of efficient taxes, we introduce the following notation. On the demand side, we let $c(q, C(\theta))$ denote the solution to (4) given consumer prices q and

[4]Pareto weights may themselves derive from political-economy considerations, like the agents' heterogeneous ability to make political contributions, as in Grossman and Helpman (1994).

aggregate consumption $C(\theta)$; we let $C(n(\theta), U(\theta))$ denote the aggregate consumption required to achieve utility $U(\theta)$ given labor supply $n(\theta)$, that is the solution to $u(C, n(\theta)) = U(\theta)$; and we let $c(q, n, U) = \int c(q, C(n(\theta), U(\theta)))dF(\theta)$ denote the total demand for goods conditional on prices, q, labor supply, $n \equiv \{n(\theta)\}$, and utility levels, $U \equiv \{U(\theta)\}$. Likewise, on the supply side, we let $y(p, n)$ denote the solution to (6) given producer prices p and labor demand n, and we let $w(p, n; \theta)$ denote the associated equilibrium wage for each household θ,

$$w(p, n; \theta) = G_{n(\theta)}(y(p, n), n) \times \frac{p \cdot y(p, n)}{\int n(\theta')G_{n(\theta')}(y(p, n), n)dF(\theta')},$$

obtained from the first-order condition with respect to $n(\theta)$, with $G_{n(\theta)} \equiv \partial G/\partial n(\theta)$.[5]

Using the previous notation, we can express the government's problem as

$$W(\phi) \equiv \max_{U,n,R\in\mathcal{R},(p,q)\in\mathcal{P}} \int U(\theta)d\Lambda(\theta) \tag{8}$$

subject to

$$n(\theta), U(\theta) \in \text{argmax}_{\tilde{n}(\theta),\tilde{U}(\theta)}\{\tilde{U}(\theta)|e(q, C(\tilde{n}(\theta), \tilde{U}(\theta))) = R(w(p, n; \theta)\tilde{n}(\theta); \theta)\}, \text{ for all } \theta,$$

$$G^*(c(q, n, U) - y(p, n); \phi) \leq 0.$$

The first constraint is the counterpart to the agent θ's upper-level optimality condition (5), whereas the second constraint is the counterpart to the good market clearing condition (1). The lower-level optimality condition (4) is already captured by the household demand, $c(q, C(n(\theta), U(\theta)))$. Likewise, the profit-maximization of old technology firms, condition (6), is already captured by the supply and wage schedules, $y(p, n)$ and $w(p, n; \theta)$. Finally, since the taxes on new technology firms t^* are unrestricted, they can always be chosen such that the profit-maximization of new technology firms, condition (7), holds for any feasible vector of output, y^*.[6]

[5]In general, both the demand of the households, $c(q, C(\theta))$, and the supply of old technology firms, $y(p, n)$, may be sets. This happens whenever households and firms operate on flat portions of their indifference curves and production possibility frontiers, respectively. Though our results generalize to environments where such situations may arise, we will ignore them for expositional purposes.

[6]This is true regardless of whether y^* is on the production possibility frontier of new technology firms. If not, taxes on new firms should just be set to $t^* = \infty$. We will ignore this knife-edge case.

4 When Is Technological Change Welcome?

We first study the welfare impact of a change in the new technology under the assumption that efficient taxes are in place.

4.1 An Envelope Result

Consider a small productivity shock from ϕ to $\phi + d\phi$. By the Envelope Theorem, we have

$$\frac{dW}{d\phi} = \gamma \frac{\partial G^*}{\partial \phi}. \tag{9}$$

This leads to our first proposition.

Proposition 1. *Technological change increases social welfare, $dW/d\phi \geq 0$, if and only if it expands the production possibility of new technology firms, that is, if and only if $\frac{\partial G^*}{\partial \phi} \leq 0$.*

Our envelope condition can be thought of as a generalization of Hulten's (1978) Theorem. In spite of the economy not being first-best, the welfare impact of technological progress can be measured in the exact same way as in first-best environments. In fact, a weaker sufficient condition for Proposition 1 to hold is the existence of a set of feasible allocation, $\mathcal{Z}$, independent of ϕ, such that the government's problem can be expressed as

$$W(\phi) = \max_{(U,c,n,y) \in \mathcal{Z}} \{ \int U(\theta) d\Lambda(\theta) | G^*(\int c(\theta) dF(\theta) - y; \phi) \leq 0 \}.$$

This formulation allows for arbitrary preferences and technology across agents and old technology firms; it also allows for production and consumption externalities as well as various market imperfections. In particular, there may be price and wage rigidities leading to labor market distortions. The first-order welfare effect of a productivity shock remains given by $\gamma(\partial G^* / \partial \phi)$.

In the presence of distortions, it is well-known that technological progress may lower welfare. This is what Edgeworth (1884) and Bhagwati (1958) refer to as "Economic Damnification" and "Immesirizing Growth". How does Proposition 1 rule it out?

The critical assumption here is not that there are no distortions or, equivalently, that our planner has enough tax instruments to target any underlying distortion. Restrictions on $\mathcal{R}$ and $\mathcal{P}$ may prevent that. Proposition 1 instead builds on the assumption that, in spite of distortions and restrictions on the set of available instruments, through $\mathcal{R}$ and $\mathcal{P}$, our planner still has enough tax instruments, through $t^* \in [-1, \infty)^N$, to control fully the behavior of new technology firms, which is where technological change is occurring.

Take the example of international trade: $G^*(y^*;\phi) = \bar{p}(\phi) \cdot y^*$. Equation (9) implies that a country gains from a terms-of-trade shock if and only if it raises the value of its net exports, evaluated at the initial quantities,

$$W'(\phi) \propto \bar{p}'(\phi) \cdot y^*.$$

Intuitively, if world prices $\bar{p}(\phi)$ change, the government always has the option to maintain domestic prices (q, p) unchanged by raising trade taxes by $\bar{p}'(\phi)$. Starting from a constrained optimum, the only first-order effect of such a policy change would be to raise tax revenues by $\bar{p}'(\phi) \cdot y^*$, regardless of whether or not the economy is first-best. Raising trade taxes by $\bar{p}'(\phi) \cdot y^*$, of course, may not be the optimal response, but the possibility of such a response is sufficient to evaluate the welfare impact of a terms-of-trade shock at the margin.[7]

4.2 Implications

We now illustrate the usefulness of our envelope result through two applications.

Taxation of Innovation. Consider first the issue of whether governments may ever want to tax innovation because of its adverse distributional consequences. To shed light on this issue in the simplest possible way, suppose that there exists a set of feasible new technologies, Φ, that can be restricted by the government. The profit maximization problem of new technology firms is now given by

$$y^*, \phi \in \text{argmax}_{\tilde{y}^*, \tilde{\phi} \in \bar{\Phi}} \{p^* \cdot \tilde{y}^* | G^*(\tilde{y}^*; \tilde{\phi}) \leq 0\},$$

where $\bar{\Phi} \subset \Phi$ is the set of technologies allowed by the government. The government's problem, in turn, takes the general form,

$$\max_{(U,c,n,y) \in \mathcal{Z}, \phi \in \Phi} \{\int U(\theta) d\Lambda(\theta) | G^*(\int c(\theta) dF(\theta) - y; \phi) \leq 0\}.$$

[7]The previous observation does not depend on whether the country is a small open economy or not. $\bar{p}(\phi)$ could itself be a function of y^*. The same result would hold. In fact, our envelope result provides the basis for a generalization of the results in Bagwell and Staiger (1999, 2001, 2012a,b): if trade taxes are unrestricted, the only rationale for a trade agreement is to correct terms-of-trade externalities.

By equation (9), the optimal technology $\bar{\phi}$ simply satisfies

$$\frac{\partial G^*(y^*;\bar{\phi})}{\partial \phi} = 0.$$

Provided that taxes of new technology firms, t^*, have been set such that they find it optimal to produce y^*, conditional on $\bar{\phi}$, they will also find it optimal to choose $\bar{\phi}$, if allowed to do so. It follows that the government does not need to affect the direction of innovation, in spite of its potential distributional implications.[8]

Valuation of New Technologies. Next we turn to the issue of evaluating the welfare gains from new technologies. Using our envelope result, one can follow, in a general equilibrium environment with distortions, the same steps used to compute equivalent and compensating variations in standard consumer theory.

Consider the following generalized version of our government's problem

$$W(\phi, D) = \max_{(U,c,n,y)\in\mathcal{Z},\phi\in\Phi} \{\int U(\theta)d\Lambda(\theta)|G^*(\int c(\theta)dF(\theta) - y;\phi) \leq D\}. \tag{10}$$

The parameters ϕ and D play the same role here as prices and income in the utility maximization problem of a single consumer. In a trade context, D corresponds to a trade deficit, that is a transfer from the rest of the world. The welfare impact of a productivity shock from ϕ_0 to ϕ_1 can then be computed either as the transfer, $EV(\phi_0, \phi_1)$, that would be equivalent to the shock, $W(\phi_0, EV(\phi_0, \phi_1)) = W(\phi_1, 0)$, or as the transfer, $CV(\phi_0, \phi_1)$, required to compensate for the shock, $W(\phi_1, -CV(\phi_0, \phi_1)) = W(\phi_0, 0)$.

Let $y^*(\phi, D)$ denote the vector of output by new technology firms associated with the solution to (10). Equation (9) implies that

$$EV(\phi, \phi + d\phi) = CV(\phi, \phi + d\phi) = \frac{\partial G^*(y^*(\phi, 0);\phi)}{\partial \phi}. \tag{11}$$

This is the counterpart to Shephard's Lemma in standard consumer theory. And, like in standard consumer theory, this envelope condition can be integrated to compute the welfare impact of arbitrary productivity shocks, as described in Appendix (A.1). Gains from trade, for instance, can be computed by integrating (11) between the current productivity level and the one at which the economy reaches autarky, as in Arkolakis, Costinot and

[8] More generally, there may be externalities across firms that directly call for subsidizing, or taxing, innovation. In such environments, the implication of our envelope result is that distributional concerns do not add a new motive for taxing, or subsidizing, innovation.

Rodríguez-Clare (2012).

Note that distortions will, in general, affect $y^*(\phi, D)$. So, the point here is not that distortions do not matter for the welfare consequences of globalization or automation. The point rather is that like in a first-best environment, the demand for goods produced using the new technology, either Chinese imports or robots, fully reveals the welfare gains associated with that technology. Concerns for redistribution and other potential sources of distortions affects how much we trade or how much we use robots, but not the mapping between quantities demanded, productivity shocks, and welfare.[9]

5 Should We Tax New Technologies?

New technologies improve efficiency, but may have adverse distributional consequences. Depending on the availability of tax instruments, governments may choose different strategies to manage these technologies. We first describe how existing benchmark results map into the government's problem of Section 3.2. We then propose to depart from them by imposing anonymous income taxation.

5.1 Negative Benchmark Results

Unrestricted Taxes. Consider first the extreme case where all tax instruments are unrestricted. This implies, in particular, that agent-specific lump-sum transfers are available, like in the Second Welfare Theorem. In this case, the government's problem can be rearranged as

$$\max_{U,n,p,q} \{ \int U(\theta) d\Lambda(\theta) | G^*(c(q,n,U) - y(p,n); \phi) \leq 0 \}.$$

The associated first-order conditions with respect to p and q are given by

$$D_p y \cdot \nabla_{y^*} G^* = 0,$$
$$D_q c \cdot \nabla_{y^*} G^* = 0,$$

[9] The extent to which our envelope result is useful in practice, of course, depends on whether constrained, but optimal policies are in place. The assumption that the government fully controls the new technology clearly is non-trivial. We view it, however, as a useful benchmark, not necessarily stronger than the opposite assumption, implicit in many papers, that governments cannot control the new technology at all. Galle, Rodríguez-Clare and Yi (2017), and Waugh and Lyon (2017) are recent welfare analysis of the so-called China shock that fall into this category. Brexit and the current debate about the renegotiation of NAFTA are stark reminders that trade policies are not set in stone.

where $D_p y \equiv \{dy_j/dp_i\}$ and $D_q c \equiv \{dc_i/dq_j\}$ are $N \times N$ matrices and $\nabla_{y^*} G^* = \{dG^*/dy_j^*\}$ is a $N \times 1$ vector.

To go from the previous first-order conditions to the optimal price gaps, note that profit maximization by new technology firms requires marginal rates of transformation to be equal to the relative prices that they face. Thus, the vector of new prices, p^*, must be collinear with $\nabla_{y^*} G^*$, and, in turn, we must have $D_p y \cdot p^* = D_q c \cdot p^* = 0$. Likewise, profit maximization by old technology requires $D_p y \cdot p = 0$—changes in output can only have second-order effects on firms' revenues—whereas expenditure minimization by agents requires $D_q c \cdot q = 0$—changes in consumption can only have second-order effects on agents' expenditure. Combining the previous observations, we get

$$D_p y \cdot (p - p^*) = 0,$$
$$D_q c \cdot (q - p^*) = 0.$$

Under the normalization, $p_1 = p_1^* = q_1 = 1$, all prices, p, p^*, and q should be equal and, accordingly, all good taxes should be zero. Not surprisingly, if lump-sum transfers are available, there is no trade-off between efficiency and redistribution.

Unrestricted Linear Taxes. Suppose now that t and t^* remain unrestricted, but that labor taxation is restricted to linear taxes,

$$R(w(\theta)n; \theta) = (1 - T(\theta))w(\theta)n \text{ for all } \theta.$$

This is the case of unrestricted linear taxation considered by Diamond and Mirrlees (1971b) and Dixit and Norman (1986). Let $r(\theta) \equiv (1 - T(\theta))w(\theta)$ denote the wage schedule faced by workers. In this case, the government's problem is

$$\max_{U,n,p,q,r} \int U(\theta) d\Lambda(\theta)$$

subject to

$$n(\theta), U(\theta) \in \operatorname{argmax}_{\tilde{n}(\theta), \tilde{U}(\theta)} \{U | e(q, C(\tilde{n}(\theta), \tilde{U}(\theta))) = r(\theta)n\}, \text{ for all } \theta,$$

$$G^*(c(q,n,U) - y(p,n); \phi) \leq 0.$$

Since p only appears in the good market clearing condition, the associated first-order condition is unchanged,

$$D_p y \cdot \nabla_{y^*} G^* = 0.$$

For the same reason as before, we therefore still have the equality of p and p^*. This is Diamond and Mirrlees's (1971b) production efficiency result: there should be no differences between the taxes faced by old and new technology firms. In a trade context, this implies that the solution to the government's problem may feature consumer taxes, $q \neq p = p^*$, but not trade taxes. A corollary of this observation is that the government can always find an allocation with consumer taxes, $t \neq 0$, and no trade taxes, $t^* = 0$ that leads to higher welfare than the autarky equilibrium, i.e, an equilibrium with prohibitive trade taxes. This is Dixit and Norman's (1986) result.

5.2 An Alternative Environment with Limited Factor Taxation

The assumption that lump-sum transfers are available or that all factors can be taxed at a different rate are clearly strong ones. In the trade policy literature, it is common to make the other extreme assumption that the only instruments available for redistribution are trade taxes. In this paper, we wish to explore further the intermediate, and more realistic, case where lump-sum transfers and factor taxation à la Diamond and Mirrlees (1971b) and Dixit and Norman (1986) are not available, but some form of income taxation may still be available and used for redistributive purposes.

In the rest of this paper, we assume that factor-specific taxes are unavailable, but that income taxation is:

$$R(w(\theta)n;\theta) = w(\theta)n - T(w(\theta)n) \text{ for all } \theta.$$

We do not impose any constraint on the shape of the retention function, R, beside anonymity. For now, we also remain agnostic about the set of feasible good prices for old technology firms and agents, $\mathcal{P}$.

In this intermediate environment, the government's problem can be expressed as

$$\max_{U,n,(p,q)\in\mathcal{P}} \int U(\theta)d\Lambda(\theta)$$

subject to

$$\theta \in \text{argmax}_{\tilde{\theta}} u\left(C(n(\tilde{\theta}), U(\tilde{\theta})), n(\tilde{\theta})\frac{w(p,n;\tilde{\theta})}{w(p,n;\tilde{\theta})}\right),$$

$$G^*(c(q,n,U) - y(p,n);\phi) \leq 0.$$

The first constraint is an incentive compatibility (IC) constraint ensuring that $\tilde{\theta} = \theta$ is optimal.[10] Conversely, for any allocation (U, n) satisfying incentive compatibility, we can find a tax schedule, T, and a retention function, R, such that the constraint on the optimality of $n(\theta)$ in the original government's problem holds for all θ.

Let $u_n \equiv \partial u / \partial n$ denote the partial derivative of u with respect to n. The envelope condition associated with the IC constraint gives

$$U'(\theta) = -u_n(C(n(\theta), U(\theta)), n(\theta))n(\theta)\omega(p, n; \theta)$$

where $\omega(p, n; \theta)$ is a local measure of wage inequality,

$$\omega(p, n; \theta) \equiv \frac{w_\theta(p, n; \theta)}{w(p, n; \theta)},$$

with $w_\theta \equiv \partial w / \partial \theta$. For piecewise differentiable allocations, the envelope condition and monotonicity of the mapping from wages, $w(p, n; \theta)$, to before-tax earnings, $w(p, n; \theta)n(\theta)$ is equivalent to incentive compatibility. We will focus on cases where $w(p, n; \theta)$ is increasing in θ, which for a given allocation can be interpreted as a normalization or ordering of θ. Under the previous conditions, we can rearrange our planning problem as

$$\max_{U, n, (p,q) \in \mathcal{P}} \int U(\theta) d\Lambda(\theta) \tag{12a}$$

subject to

$$U'(\theta) = -u_n(C(n(\theta), U(\theta)), n(\theta))n(\theta)\omega(p, n; \theta), \tag{12b}$$

$$G^*(c(q, n, U) - y(p, n); \phi) \leq 0. \tag{12c}$$

5.3 A First Set of Tax Formulas

We study separately the cases with and without taxes on old technology firms.

[10] In order to achieve the earnings $n(\tilde{\theta})w(p, n; \tilde{\theta})$ of an agent of type $\tilde{\theta}$, an agent of type θ must supply $n(\tilde{\theta})w(p, n; \tilde{\theta})/w(p, n; \theta)$ units of labor. Since all agents have the same weakly separable preferences, if agent θ mimics agent $\tilde{\theta}$ and receives the same earnings, she must also achieve the same aggregate consumption, $C(n(\tilde{\theta}), U(\tilde{\theta}))$.

Case 1: Unrestricted taxes on old and new technology firms. In this case, the set of feasible prices that households and old technology firms may face is given by

$$\mathcal{P} = \{(p, q) \in \mathbb{R}_+^N \times \mathbb{R}_+^N\}.$$

The first-order conditions with respect to p and q therefore imply

$$D_p y \cdot \nabla_{y^*} G^* = -[\int \mu(\theta) u_n(\theta) n(\theta) (\nabla_p \omega(\theta)) d\theta] / \gamma,$$
$$D_q c \cdot \nabla_{y^*} G^* = 0,$$

where $\mu(\theta)$ denotes the Lagrange multiplier associated with agent θ's IC constraint, γ denotes the Lagrange multiplier associated with the good market clearing condition, $u_n(\theta) \equiv u_n(C(n(\theta), U(\theta)), n(\theta))$, and $\nabla_p \omega(\theta) \equiv \{d\omega(p, n; \theta)/dp_j\}$. Like in Section 5.1, we can rearrange the previous expressions in terms of price gaps using the fact that firms maximize profits and agents minimize expenditure,

$$D_p y \cdot (p^* - p) = -[\int \mu(\theta) u_n(\theta) n(\theta) (\nabla_p \omega(\theta)) d\theta] / [\gamma G^*_{y_1^*}], \tag{13}$$
$$D_q c \cdot (p^* - q) = 0. \tag{14}$$

For the same reason as in Section 5.1, equation (14) immediately implies that $p^* = q$ and, in turn, that there should be no taxes on new technology firms: $t^* = 0$. The implications of equation (13) for optimal taxation are more subtle. It includes the Lagrange multipliers associated with the incentive compatibility constraint, $\mu(\theta)$, and the good market clearing condition, γ. Those are neither primitives nor observables. As a first step towards operationalizing the previous results, we therefore propose to solve for $\mu(\theta)/\gamma G^*_{y_1^*}$, as a function of the distribution of Pareto weights, Λ, which we take as primitives, and observables, such as marginal income tax rates, $\tau(\theta) \equiv T'(w(\theta)n(\theta))$, and earnings, $x(\theta) \equiv w(\theta)n(\theta)$.

To do so, we can use the first-order condition with respect to $\{U(\theta)\}$ and the fact that $u_n(\theta)/u_C(\theta) = -(1 - \tau(\theta))w(\theta)$. As formally established in Appendix A, this leads to

$$D_p y \cdot (p^* - p) = \int \psi(\theta)(1 - \tau(\theta)) x(\theta) (\nabla_p \omega(\theta)) d\theta, \tag{15}$$

where the weight $\psi(\theta)$ is such that

$$\psi(\theta) = \int_{\theta}^{\bar{\theta}} \exp[-\int_{\theta}^{z} \rho(v)dx(v)]u_C(z)\zeta(z)dF(z) \tag{16}$$
$$-\frac{\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_{\underline{\theta}}^{z} \rho(v)dx(v)]u_C(z)\zeta(z)dF(z)}{\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_{\underline{\theta}}^{z} \rho(v)dx(v)]u_C(z)d\Lambda(z)} \int_{\theta}^{\bar{\theta}} \exp[-\int_{\theta}^{z} \rho(v)dx(v)]u_C(z)d\Lambda(z),$$

with $u_C(\theta) \equiv u_C(C(n(\theta), U(\theta)), n(\theta))$ the marginal utility of aggregate consumption of agent θ, $\rho(\theta) \equiv w(\theta)\frac{\partial(u_n/u_C)}{\partial C}$ the partial derivative, with respect to aggregate consumption, of the marginal rate of substitution between earnings and aggregate consumption, and $\zeta(\theta) \equiv p^* \cdot c_{U(\theta)}(q, C(n(\theta), U(\theta)))$ the social marginal value of agent θ's utility.[11]

Noting that $p_i^* - q_i = -t_i^* p_i^*$ and $p_i - q_i = -t_i p_i$, we can summarize the implications of equations (14) and (15) for the optimal ad-valorem taxes on old and new technology firms as follows.

Proposition 2. *Suppose that taxes on both new and old technology firms are unrestricted. Then Pareto efficient taxes satisfy*

$$t^* = 0,$$

$$D_p y \cdot (tp) = \int \psi(\theta)(1 - \tau(\theta))x(\theta)(\nabla_p \omega(\theta))d\theta,$$

where $tp \equiv \{t_i q_i\}$ is a $N \times 1$ vector and $\psi(\theta)$ is given by equation (16).

The fact that optimal taxes on new technology firms are zero is an expression of the targeting principle. Here, as in Naito (1999), the rationale for good taxation is to manipulate wages and this is best achieved by manipulating the prices of old technology firms that, together with labor supply, determine these wages.[12] Specifically, by affecting the prices faced by old technology firms, the government may lower inequality, as measured by a decrease in $\omega(\theta)$, and relax the incentive compatibility constraint of agent θ. The lower inequality is, the more costly it becomes for more skilled agents to mimic the behavior of less skilled agents, and hence the lower their informational rents. This creates

[11] When $p^* = q$, as implied by equation (14), $\zeta(\theta)$ reduces to the inverse of agent θ's marginal utility of income, $q \cdot c_{U(\theta)}(q, C(n(\theta), U(\theta)))$. When $p^* \neq q$, as will be the case below, $\zeta(\theta)$ also includes the change in fiscal revenue, $(p^* - q) \cdot c_{U(\theta)}(q, C(n(\theta), U(\theta)))$.

[12] Mayer and Riezman (1987) establish a similar result in a trade context with inelastic factor supply and no income taxation. If both producer and consumer taxes are available, they show that only the former should be used. This result, however, requires preferences to be homothetic, as discussed in Mayer and Riezman (1989). Our result does not require this restriction. This reflects the fact that we have access to non-linear income taxation, leading to the envelope condition (12b) rather than the budget constraint, $e(q, u(\theta)) = w(p, n, \theta)n(\theta)$, in our planning problem. We come back to this point in footnote 13.

a first-order welfare gains from taxes on old technology firms. Intuitively, it is cheaper to achieve redistribution through income taxation if earnings, before income taxes, are already more equal.

The key difference between Proposition 2 and the work of Naito (1999), beside greater generality, is the fact that our analysis goes beyond the first-order condition (13) by solving for the endogenous Lagrange multipliers, $\mu(\theta)/\gamma G^*_{y_1^*}$, as a function of primitives and observables. Under the assumption that preferences are quasi-linear, $U(\theta) \equiv c_1(\theta) + \sum_{i\neq 1} v_i(c_i(\theta)) - h(n(\theta))$, this provides a particularly simple optimal tax formula. In this case, since $u_C(\theta) = 1$, $\zeta(\theta) = 1$, and $\rho(\theta) = 0$ for all θ, we have

$$\psi(\theta) = \Lambda(\theta) - F(\theta). \tag{17}$$

This leads to the following corollary.

Corollary 1. *Suppose that taxes on both new and old technology firms are unrestricted and that upper-level preferences are quasi-linear. Then Pareto efficient taxes on old technology firms satisfy*

$$D_p y \cdot (tp) = \int (\Lambda(\theta) - F(\theta))(1 - \tau(\theta))x(\theta)(\nabla_p \omega(\theta))d\theta.$$

Case 2: No taxes on old technology firms. We now briefly discuss the case where only taxes on new technology firms are available. Formally, this corresponds to assuming

$$\mathcal{P} = \{(p, q) \in \mathbb{R}^N_+ \times \mathbb{R}^N_+ | p = q\}.$$

Under this additional constraint, the first-order condition with respect to p becomes

$$(D_p y - D_q c) \cdot (p^* - p) = -\int \frac{\mu(\theta)}{\gamma} u_n(\theta) n(\theta)(\nabla_p \omega(\theta))d\theta,$$

The rest of the analysis is unchanged. In particular, equation (16) still holds. These two observations lead to our next proposition.

Proposition 3. *Suppose that only taxes on new technology firms are available. Then Pareto efficient taxes satisfy*

$$(D_q c - D_p y) \cdot (p^* t^*) = \int \psi(\theta)(1 - \tau(\theta))x(\theta)(\nabla_p \omega(\theta))d\theta,$$

where $p^* t^* \equiv \{p_i t_i^*\}$ is a $N \times 1$ vector and $\psi(\theta)$ is given by equation (16).

The economics is the same as in the case of Proposition 2. The only difference is that a change in p now distorts both consumption and output decisions. Hence, it is the change in output net of consumption, $D_p y - D_q c$, that matters for the optimal level of the taxes.

In the case of quasi-linear upper-level preferences, we still have $\psi(\theta) = \Lambda(\theta) - F(\theta)$, leading to the counterpart of Corollary 1.

Corollary 2. *Suppose that only taxes on new technology firms are available and that upper-level preferences are quasi-linear. Then Pareto efficient taxes satisfy*

$$(D_q c - D_p y) \cdot (p^* t^*) = \int (\Lambda(\theta) - F(\theta))(1 - \tau(\theta))x(\theta)(\nabla_p \omega(\theta))d\theta.$$

In a trade context, Corollary 2 implies that two key determinants of optimal tariffs are: (i) the difference between the Pareto weights of the government, Λ, and the utilitarian one, F; and (ii) the elasticity of import demand, as captured by $D_q c - D_p y$. These are the same determinants found in the optimal tariff formula of Grossman and Helpman (1994), where $\Lambda(\theta)$ reflects whether agents are politically organized or not.

This should be intuitive. The key difference between the class of problems that we consider and those in the political economy of trade literature is that we allow for endogenous labor supply and income taxation. So far, however, we have not used the first-order condition with respect to n, which explains the similarity between Proposition 3 and the existing trade literature.[13] Here, labor supply considerations are implicitly captured by the optimal marginal tax rates, $\tau(\theta)$. We explore them in detail in the next section.

[13]Formally, in the absence of income taxation and with inelastic factor supply, the government's problem of Section 3.2 reduces to

$$\max_{U,(p,q)\in\mathcal{P}} \int U(\theta) d\Lambda(\theta),$$

subject to

$$e(q, U(\theta)) = w(p, n; \theta)n(\theta),$$
$$G^*(c(q, n, U) - y(p, n); \phi) \leq 0.$$

Under the restrictions that $p = q$, the first-order condition with respect to p is given by

$$(D_q c - D_p y) \cdot \nabla G^*_{y^*} = \int \frac{\mu(\theta)}{\gamma} [\nabla_p e - (\nabla_p w)n(\theta)] d\theta \neq 0.$$

Taking the first-order condition with respect to $U(\theta)$, as we did above, one can further relate $\mu(\theta)/\gamma$ to the Pareto weights to obtain Grossman and Helpman's (1994) formula.

5.4 Correlations and Bounds

Correlations. Since the rationale behind taxes on old and new technology firms is to redistribute from the rich to the poor, by lowering inequality and relaxing IC constraints, it seems natural to expect, on average, higher taxes on old technology firms, relative to new technology firms, on goods for which higher prices, p_i, are associated with more inequality. Our next result formalizes this intuition.

Since old technology firms maximize profits, $D_p y$ must be positive semi-definite. Likewise, since agents minimize expenditure, $D_q c$ must be negative semi-definite. For any vector of price gaps, $p^* - p$, we must therefore have

$$(p^* - p)' \cdot D_p y \cdot (p^* - p) \geq 0,$$
$$(p^* - p)' \cdot D_q c \cdot (p^* - p) \leq 0,$$

where $(p^* - p)'$ denotes the transpose of $(p^* - p)$. Using these two observations, we obtain the following corollary of Propositions 2 and 3.

Corollary 3. *Regardless of whether or not only taxes on new technology firms are available, optimal price gaps between new and old technology firms satisfy*

$$(p^* - p)' \cdot \int \psi(\theta)(1 - \tau(\theta))x(\theta)(\nabla_p \omega(\theta))d\theta \geq 0.$$

In words, old technology firms should tend to have lower prices, i.e., be taxed more relative to new technology firms, in sectors that tend to increase inequality the most, i.e., those for which $\int \psi(\theta)(1 - \tau(\theta))x(\theta)(\omega_{p_i}(\theta))d\theta$ is high. If taxes on new and old technology firms are available, this will take the form of higher taxes, t_i, on old technology firms. If only the former taxes are available, this will take the form of lower taxes, t_i^*, on new technology firms.

Bounds. Up to this point, all our results about the structure of optimal taxes requires knowledge of the government's Pareto weights, Λ. We conclude this section by providing bounds on optimal taxes that dispense with such information. We focus on the case with quasi-linear upper-level preferences discussed in Corollaries 1 and 2. This is the case that we will study further in the next section.[14]

[14]Similar bounds can be obtained in the case with weakly separable, but not necessarily additively separable preferences.

From equation (17) and the fact that $\Lambda(\theta) \in [0,1]$, we know that

$$-F(\theta) \leq \psi(\theta) \leq 1 - F(\theta).$$

Let $\Theta_i^+ \equiv \{\theta \in [\underline{\theta}, \bar{\theta}] | \omega_{p_i}(\theta) > 0\}$ and $\Theta_i^- \equiv \{\theta \in [\underline{\theta}, \bar{\theta}] | \omega_{p_i}(\theta) < 0\}$ denote the set of agents for which an increase in p_i locally raises and lowers inequality, respectively. Using the previous notation, we obtain the following corollary of Proposition 2.

Corollary 4. *Suppose that taxes on both new and old technology firms are unrestricted. Then Pareto efficient taxes on old technology firms satisfy*

$$D_{p_i}y \cdot (tp) \leq \int (\mathbf{1}_{\Theta_i^+}(\theta) - F(\theta))(1 - \tau(\theta))x(\theta)\omega_{p_i}(\theta)d\theta,$$
$$D_{p_i}y \cdot (tp) \geq \int (\mathbf{1}_{\Theta_i^-}(\theta) - F(\theta))(1 - \tau(\theta))x(\theta)\omega_{p_i}(\theta)d\theta,$$

where $D_{p_i}y$ denotes the i-th row of D_py and $\mathbf{1}_{\Theta}(\cdot)$ is the indicator function of $\Theta \in \{\Theta_i^+, \Theta_i^-\}$.

Beside the distortionary impact of good taxation on output, measured by $D_{p_i}y$, Corollary 4 implies that the only information required to bound the optimal taxes on old technology firms are data on, or estimates of, the impact of good prices on inequality, $\omega_{p_i}(\theta)$, earnings, $x(\theta)$, marginal income tax rates, $\tau(\theta)$, and the distribution of skills, $F(\theta)$.[15] Note also that for any good i that raises inequality for all agents, $\omega_{p_i}(\theta) > 0$ for all θ, the upper-bound is simply given by $\int(1 - F(\theta))(1 - \tau(\theta))x(\theta)\omega_{p_i}(\theta)dF(\theta)$, which corresponds to optimal price gap of a government with Rawlsian preferences ($\Lambda(\theta) = 1$ for all θ).

Starting from Proposition 3, the same arguments can be used to provide bounds on the optimal taxes on new technology firms, when only those taxes are available.

6 Applications: Robots and Trade

To provide further intuition about the forces that shape Pareto efficient taxes, we turn to a special case of the economic environment presented in Section 2. There is one final good, indexed by f, and one intermediate good, indexed by m, which could be either robots that are produced domestically or machines that are imported from abroad. The questions that we are interested in are: Can one characterize Pareto efficient taxes on robots or trade without making arbitrary assumptions on the distributions of Pareto weights? Given the estimated impact of robots or trade on wage inequality, how large should Pareto efficient

[15] Since we can always change variable and express the above integrals as a function of earnings, this is equivalent to having access to data on the distribution of earnings.

taxes be? If robots or traded goods were to become cheaper, and exacerbate inequality further, should taxes be raised further as well?

6.1 A Two-Good Economy

Agents have quasi-linear preferences

$$U(\theta) = C(\theta) - h(n(\theta)). \tag{18}$$

where $C(\theta)$ now denotes the consumption of the unique final good, which we use as our numeraire, $q_f = p_f = p_f^* = 1$. Old technology firms produce the final good, $y_f \geq 0$, using workers, $n \equiv \{n(\theta)\}$, and machines, $y_m \leq 0$, as an input. The production set is given by

$$G(y_f, y_m, n) = y - \max_{\{\tilde{y}_m(\theta)\}} \{\int g(\tilde{y}_m(\theta), n(\theta); \theta) dF(\theta) | \int \tilde{y}_m(\theta) dF(\theta) \leq -y_m\}, \tag{19}$$

where $g(\cdot, \cdot; \theta)$ is homogeneous of degree one for all θ and $\tilde{y}_m(\theta)$ represents the number of machines combined with workers of type θ to produce the final good. Like in the examples of Section 2.2, new technology firms produce machines, $y_m^* \geq 0$, using the final good, $y_f^* \leq 0$,

$$G^*(y_f^*, y_m^*) = \phi y_f^* + y_m^*, \tag{20}$$

where ϕ measures the productivity of machine producers.

Let p_m and p_m^* denote the price of robots faced by old and technology firms. Since machines are only demanded by firms, but not households, these are the only relevant good prices in this economy. Profit maximization by new technology firms implies

$$p_m^* = 1/\phi,$$

whereas profit maximization by old technology firms implies

$$1 = b(p_m, w(\theta); \theta),$$

where $b(p_m, w(\theta); \theta) \equiv \min_{\tilde{y}_m(\theta), \tilde{n}(\theta)} \{p_m \tilde{y}_m(\theta) + w(\theta) n | g(\tilde{y}_m(\theta), \tilde{n}(\theta); \theta) \geq 1\}$ denotes the unit costs of firms using agents of type θ. Note that the previous expression implies that the wage $w(\theta)$ of any agent θ only depends on the price of machines p_m faced by old technology firms, not on labor supply decisions, $n \equiv \{n(\theta)\}$. Here, because of additive separability in production, machines directly affect inequality by affecting relative marginal products of labor, but not indirectly through further changes in relative labor supply.[16]

In line with the notation of the previous sections, we let $w(p_m;\theta)$ denote the equilibrium wage of agent θ as a function of the price of robots, that is the unique solution to $1 = b(p_m, w(\theta);\theta)$. Similarly, we let $y_m(p_m, n(\theta);\theta)$ denote the equilibrium number of machines used by agents of type θ, that is the solution to $p_m = dg(y_m(\theta), n(\theta);\theta)/dy_m(\theta)$; we let $y_m(p_m, n) \equiv -\int y_m(p_m, n(\theta);\theta)dF(\theta)$ denote the total demand for machines by old technology firms; and we let $y_f(p_m, n) \equiv \int g(y_m(p_m, n(\theta);\theta)), n(\theta);\theta)dF(\theta)$ denote the total supply of the final good by these firms.

Under the previous assumptions, the planner's problem (12) simplifies into

$$\max_{U,n,p_m} \int U(\theta)d\Lambda(\theta)$$

subject to

$$U'(\theta) = h'(n(\theta))n(\theta)\omega(p_m;\theta),$$

$$\phi[\int (U(\theta) + h(n(\theta)))dF(\theta) - y_f(p_m, n)] - y_m(p_m, n) \leq 0.$$

Here, the final constraint states that total consumption of the final good must be weakly less than gross output, $y_f(p_m, n)$, minus investment in machines, $-y_m(p_m, n)/\phi$.

6.2 Another Tax Formula

To characterize the optimal tax on robots, we again start from the first-order condition with respect to p_m. In this environment, equation (13) simplifies into

$$\gamma\phi(p_m^* - p_m)y_{m,p_m} = \int \mu(\theta)h'(n(\theta))n(\theta)\omega_{p_m}(p_m;\theta)d\theta, \tag{21}$$

with $y_{m,p_m} \equiv \frac{dy_m(p_m,n)}{dp_m}$. The first-order condition with respect to $n(\theta)$ also takes a simple form. As demonstrated in Appendix A.3, it can be expressed as

$$\gamma\phi[w(\theta)\tau(\theta) + (p_m^* - p_m)y_{m,n(\theta)}]f(\theta) = \mu(\theta)h'(n(\theta))[\frac{\epsilon(\theta)+1}{\epsilon(\theta)}]\omega(p_m;\theta), \tag{22}$$

where $\epsilon(\theta) \equiv \frac{d\ln(n(\theta))}{d\ln h'(n(\theta)))} \geq 0$ denotes the labor supply elasticity and $y_{m,n(\theta)} \equiv \frac{dy_m(p_m,n)}{dn(\theta)}$. Using the previous expression to substitute for $\mu(\theta)h'(n(\theta))/\gamma$ in equation (21) and not-

[16] This second effect is the focus of the analysis of optimal income taxes in Stiglitz (1982).

ing that $d \ln y_m(p_m, n(\theta); \theta)/d \ln n(\theta) = 1$,[17] we obtain

$$p_m - p_m^* = \frac{\int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d \ln \omega(\theta)}{d \ln |y_m|} \cdot \tau(\theta) \cdot x(\theta) dF(\theta)}{|y_m|[1 - \int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d \ln \omega(\theta)}{d \ln |y_m|} \cdot \frac{y_m(\theta)}{|y_m|} dF(\theta)]}, \tag{23}$$

where $\frac{d \ln \omega(\theta)}{d \ln |y_m|}$ corresponds to the local elasticity of relative wages with respect to the total number of machines used by old technology firms, holding workers' labor supply fixed,

$$\frac{d \ln \omega(\theta)}{d \ln |y_m|} = \frac{\partial \ln p_m}{\partial \ln |y_m(p_m, n)|} \frac{d \ln \omega(p_m; \theta)}{d \ln p_m}.$$

Since households do not consume machines, the previous price gap can be implemented equivalently with a negative tax on old technology firms, a positive tax on new technology firms, or a combination of both. For expositional purposes, we focus in the rest of this section on the case where the tax on old technology firms has been set to zero and refer to $t_m^* = p_m/p_m^* - 1$ as the tax on machines. Given this normalization, equation (23) leads to the following proposition.

Proposition 4. *Suppose that equations (18)-(20) hold. Then the Pareto efficient tax satisfies*

$$t_m^* = \frac{\int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \frac{d \ln \omega(\theta)}{d \ln |y_m|} \tau(\theta) x(\theta) dF(\theta)}{p_m^* |y_m|[1 - \int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \frac{d \ln \omega(\theta)}{d \ln |y_m|} \frac{y_m(\theta)}{|y_m|} dF(\theta)]}.$$

Like the formulas of Section 5, this new tax formula gives center stage to the impact of machines on inequality, here captured by $\frac{d \ln \omega(\theta)}{d \ln |y_m|} > 0$. Unlike the formulas of Section 5, however, this new tax formula does not require any assumption about the distribution of the Pareto weights, Λ. Intuitively, the underlying preferences of the government for the utility of different groups of agents get revealed by the marginal tax rate $\tau(\theta)$ that they face after controlling for the distortionary cost of redistribution, as measured by the labor supply elasticity $\epsilon(\theta)$.[18]

An alternative perspective on Proposition 4 is that it compares the magnitude of two fiscal externalities. Consider a tax reform that simultaneously raises the tax on machines, t_m^*, and changes the income tax schedule to maintain all households at their original utility level. Such a reform would lower the total demand for machines, both directly,

[17]Recall that $y_m(p_m, n(\theta); \theta)$ is implicitly defined as the solution to $p_m = dg(y_m(\theta), n(\theta); \theta)/dy_m(\theta)$. Since $g(\cdot, \cdot; \theta)$ is homogeneous of degree one, this is equivalent to $p_m = dg(y_m(\theta)/n(\theta), 1; \theta)/dy_m(\theta)$. Differentiating, we therefore get $d \ln y_m(p_m, n(\theta); \theta)/d \ln n(\theta) = 1$.

[18]This is the idea behind Werning's (2007) test of whether an income tax schedule is Pareto optimal. Namely, it is if the inferred Pareto weights are all positive.

through the increase in the price faced by old technology firms, and indirectly, through the change in the labor supply of households. A smaller number of machines would lower the tax revenues collected on those machines in proportion to $p_m - p_m^*$, as captured by $(p_m - p_m^*)\,|y_m|[1 - \int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d\ln\omega(\theta)}{d\ln|y_m|} \cdot \frac{y_m(\theta)}{|y_m|} dF(\theta)]$. Such a reform would also affect the total taxes collected on labor earnings, which is the effect captured by $\int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d\ln\omega(\theta)}{d\ln|y_m|} \cdot \tau(\theta) \cdot x(\theta) dF(\theta)$. For the ad-valorem tax on machines t_m^* to be Pareto efficient, the marginal cost of the previous reform, in terms the fiscal losses associated with machines, should be equal to its marginal benefit, in terms of the fiscal surplus associated with labor earnings.

6.3 Putting the Formula to Work

The key input into the formula of Proposition 4 is the local elasticity of relative wages, with respect to the number of machines, across the wage distribution. We now discuss, using two examples, how estimates of this elasticity can be obtained from existing empirical work as well as the associated policy implications.

Robot Example: Acemoglu and Restrepo (2017). Our first example focuses on robots. Using a difference-in-difference strategy, Acemoglu and Restrepo (2017) have estimated the effect of industrial robots, defined as "an automatically controlled, reprogrammable, and multipurpose [machine]" on different quantiles of the wage distribution between 1990 and 2007 across US commuting zones. Using our notation, their estimates can be interpreted as the semi-elasticity of wages with respect to robots, $\eta_{AR}(\theta) = \frac{d\ln w(\theta)}{d|y_m|}$, where $|y_m|$ is expressed as number of robots per thousand workers. The elasticity that we are interested in can then be approximated by

$$\frac{d\ln\omega(\theta)}{d\ln|y_m|} \simeq \frac{|y_m|}{\Delta\ln w(\theta)} \times \Delta\eta_{AR}(\theta),$$

where $\Delta z(\theta)$ denotes the change between a given variable $z(\theta)$ between two consecutive quantiles of the wage distribution. Figure 1a reports the local elasticity of relative wages (solid line) using $|y_m| \simeq 1.2$ as the number of robots per thousand workers in the United States in 2007.

We first implement our formula abstracting away from any heterogeneity in labor supply and relative wage elasticities, $\epsilon(\theta) = \epsilon$ and $\frac{d\ln\omega(\theta)}{d\ln|y_m|} = \frac{d\ln\omega}{d\ln|y_m|}$.[19] In this case, we

[19] In the next subsection, we provide micro-foundations on preferences and technology such that both elasticities are indeed constant across agents of different types.

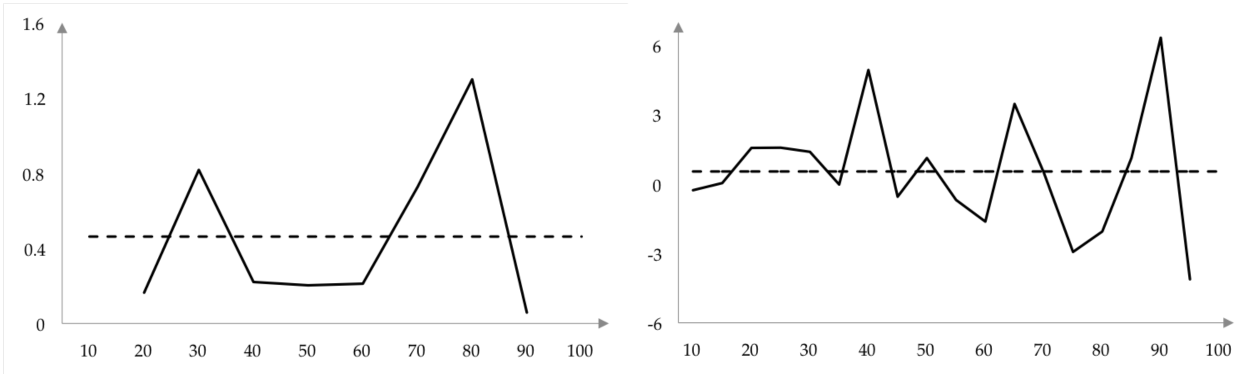

(a) Robots (Acemoglu and Restrepo, 2017) **(b)** Chinese Imports (Chetverikov, Larsen and Palmer, 2016)

Figure 1: Elasticity of relative wages, $\frac{d\ln\omega(\theta)}{d\ln y_m}$, across quantiles of US wage distribution.

can further simplify the formula of Proposition 4 into

$$t_m^* = \frac{\int x(\theta)dF(\theta)}{p_m^*|y_m|}\frac{\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}\bar{\tau}}{1-\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}},$$

where $\bar{\tau} \equiv \int \tau(\theta)\frac{x(\theta)}{\int x(\theta')dF(\theta')}dF(\theta)$ denotes the average marginal income tax rates. In Figure 1a, the average of the local elasticity, $\frac{d\ln\omega}{d\ln|y_m|}$, across deciles (dashed line) is around 0.5. Graetz and Michaels (2018) estimate that the share of robot services in total capital services is around 0.64%, which leads to a ratio of total labor earnings to spending on robots around 245, whereas Guner, Kaygusuz and Ventura (2014) point toward an average marginal income tax rate around 0.1 for the United States. For a labor supply elasticity of $\epsilon = 0.1$, in the low range of the micro-estimates reviewed in Chetty (2012), we therefore obtain an optimal tax on robots equal to $t_m^* = 99\%$. For higher labor supply elasticities of $\epsilon = 0.3$ or $\epsilon = 0.5$, the previous tax goes up to $t_m^* = 269\%$ and $t_m^* = 410\%$, respectively.

More generally, in order to take into account the heterogeneity in relative wage elasticities reported in Figure 1a when implementing the formula of 4, one would need data on the share of robots employed with workers of a particular type, $\frac{y_m(\theta)}{|y_m|}$. Since we do not have access to such data, we propose instead to compute the bounds on t_m^* associated with the minimum and maximum values of $\frac{\epsilon(\theta)}{\epsilon(\theta)+1}\frac{d\ln\omega(\theta)}{d\ln|y_m|}$. For a constant labor supply elasticity equal to $\epsilon = 0.1$, this leads to a lower-bound on the tax on robots equal to 104% and 117%, respectively.

Trade Example. Our second example focuses on Chinese imports. Using the same difference-in-difference strategy as in Autor, Dorn and Hanson (2013), Chetverikov, Larsen and Palmer (2016) have estimated the effect on log wages of a \$1,000 increase in Chinese imports per worker at different percentiles of the wage distribution. Following the same approach as in the case of robots, we can transform the previous semi-elasticity into an elasticity. Figure 1b reports the local elasticity of relative wages (solid line) using $|y_m| \simeq 2.2$ as the value of Chinese imports, in thousands of US dollars, per worker for the United States in 2007.

The average value of the relative wage elasticity, $\frac{d \ln \omega}{d \ln |y_m|}$, is of the same order of magnitude as the one implied by the estimates of Acemoglu and Restrepo (2017), around 0.5 (dashed line in Figure 1b). Compared to the robot example, however, the ratio of total labor earnings to total Chinese imports in 2007 is only 26.4, an order of magnitude smaller than the ratio to total spending on robots. Implementing again the formula of Proposition 4 in the absence of heterogeneity in labor supply and relative wage elasticities, we obtain a much smaller tariff: $t_m^* = 15\%$ for $\epsilon = 0.1$, 42% for $\epsilon = 0.3$, and 65% for $\epsilon = 0.5$.

Another key difference between the two examples is that the relative wage elasticity, $\frac{d \ln \omega(\theta)}{d \ln |y_m|}$, is much more heterogeneous across quantiles in the trade case, as can be seen in Figure 1. This leads to much broader bounds on the efficient tariffs when that heterogeneity is taken into account. For a constant labor supply elasticity $\epsilon = 0.1$, the lower-bound on the Pareto efficient tariff now is equal to 2%, whereas the upper-bound is equal to 7%.

6.4 Comparative Statics

Our final set of results explores the relationship between the efficiency of the new technology, ϕ, and the magnitude of the Pareto efficient tax.

We take a first pass at this question in the context of a parametric example with Rawlsian preferences, constant Frisch elasticities, and Cobb-Douglas production functions. Formally, we assume that the distribution of Pareto weights is such that $\Lambda(\theta) = 1$ for all θ; the distribution of types is uniformly distributed between 0 and 1; preferences and technology are such that

$$h(n(\theta)) = \frac{(n(\theta))^{1+1/\epsilon}}{1 + 1/\epsilon}, \tag{24}$$

$$g(y_m(\theta), n(\theta); \theta) = \exp(\alpha(\theta)) \cdot (\frac{y_m(\theta)}{\beta(\theta)})^{\beta(\theta)} (\frac{n(\theta)}{1 - \beta(\theta)})^{1-\beta(\theta)}, \tag{25}$$

with $\alpha(\theta) \equiv \frac{\alpha \ln(1-\theta)}{\beta \ln(1-\theta) - 1}$, $\beta(\theta) \equiv \frac{\beta \ln(1-\theta)}{\beta \ln(1-\theta) - 1}$, and $\alpha, \beta > 0$.

In this case, the zero-profit condition of old technology firms leads to

$$w(p_m;\theta) = (1-\theta)^{-1/\gamma(p_m)},$$

with $\gamma(p_m) \equiv 1/(\alpha - \beta \ln p_m)$. Under the restriction that $\gamma(p_m) > 0$, which we maintain throughout, wages are increasing in θ and Pareto distributed with shape parameter equal to $\gamma(p_m)$ and lower bound equal to 1. By construction, more skilled workers tend to use machines relatively more; $\beta(\theta)$ is increasing in θ. So an increase in the price of machines tends to lower their wages relatively more and decrease inequality,

$$\frac{d\ln\omega(\theta)}{d\ln p_m} = -\frac{d\ln\gamma(p_m)}{d\ln p_m} = -\beta\gamma(p_m) < 0.$$

For comparative static purposes, a limitation of the formula in Proposition 4 is that it involves the entire schedule of optimal marginal tax rates $\{\tau(\theta)\}$. These are themselves endogenous objects that will respond to productivity shocks. In Appendix A.4, we demonstrate how to substitute for those by combining the first-order conditions with respect to $U(\theta)$ and $n(\theta)$. This leads to

$$\frac{t_m^*}{1+t_m^*} = \frac{\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}\tau^*}{1-\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}\tau^*}\frac{1-s_m}{s_m}, \tag{26}$$

where the elasticity of relative wages, $\frac{d\ln\omega}{d\ln|y_m|} \equiv -\beta\gamma(p_m)\frac{\partial\ln p_m}{\partial\ln|y_m(p_m,n)|}$, is now constant across agents; $\tau^* \equiv \frac{\epsilon+1}{\epsilon+1+\epsilon\gamma(p_m)}$ corresponds to the optimal marginal tax rate that would be imposed in the absence of a tax on machines, as in Diamond (1998), Saez (2001), and Scheuer and Werning (2017); and $s_m \equiv \frac{p_m|y_m|}{\int x(\theta)dF(\theta)+p_m|y_m|}$ measures the share of machines in gross output.

After expressing the three previous statistics, $\frac{d\ln\omega}{d\ln|y_m|}$, τ^*, and s_m, as functions of t_m^* and ϕ, we can apply the Implicit Function Theorem to determine the monotonicity of the Pareto efficient tax, as we do in Appendix A.4. This leads to our final proposition.

Proposition 5. *Suppose that equations (18)-(25) hold. Then the Pareto efficient tax, t_m^*, decreases with the productivity of machine makers, ϕ.*

In this economy, more machines always increase inequality, $\frac{d\ln\omega}{d\ln|y_m|} > 0$, but for comparative static purposes, the relevant question is whether this effect gets exacerbated as the new technology improves. Here, one can check that $\frac{\partial}{\partial\phi}\frac{d\ln\omega}{d\ln|y_m|} < 0$ both because relative wages are becoming less responsive to the price of machines, $\frac{\partial}{\partial\phi}|\frac{d\ln\omega}{d\ln|p_m|}| < 0$, and because the demand for machines is becoming more elastic, $\frac{\partial}{\partial\phi}|\frac{\partial\ln p_m}{\partial\ln|y_m(p_m,n)|}| < 0$, due to the increase in the labor supply of high-skilled workers whose demand for machines is more elastic. One can also check that these two effects dominate the increase in the marginal tax rate, $\frac{\partial \tau^*}{\partial \phi} > 0$, in response to greater inequality. For a given share of machines s_m, this implies that the total fiscal externality associated with new machines decreases. Since the share of machines increase with improvements in the new technology, $\frac{\partial s_m}{\partial \phi} > 0$, the fiscal externality per machine a fortiori decreases and so does the tax on machines.

As this example illustrates, cheaper robots may lead to a higher share of robots in the economy, more inequality, but a lower optimal tax on robots. Likewise, more imports and more inequality, in spite of the government having extreme distributional concerns and imports causing inequality, may be optimally met with less trade protection. This does not occur because redistribution is becoming more costly as the economy gets more open.[20] Here, the elasticity of labor supply is fixed and τ^* increases with ϕ. This also does not occur because redistribution through income taxation is becoming more attractive. Everything else being equal, an increase in τ^* raises the tax on imports. Rather the decrease in the tax on imports captures a standard Pigouvian intuition: as ϕ increases, the total fiscal externality associated with imports increases, but the externality per unit of imports does not, leading to a lower value of the Pareto efficient tax.

7 Conclusion

Our paper has focused on two broad sets of issues. The first one is related to the evaluation of new technologies. In spite of tax instruments being limited and the government having concerns for redistribution, we have shown that productivity shocks can be evaluated using a simple envelope argument, like in first best environments. Our finding stands in sharp contrast with the existing results of a large literature concerned with distortions and welfare. Our envelope result implies that firms developing new technologies should not be taxed or subsidized. It also implies that the welfare gains from trade or the welfare gains from the introduction of robots can still be computed by integrating below the demand curves for foreign goods and robots, respectively. Distributional concerns and various distortions may affect how much we trade or how much we use robots, but not the welfare implications from changes in the demand for those.

The second set of issues is related to the management of new technologies. We have asked: should we tax or subsidize firms using new technologies? And to the extent that

[20] This is the point emphasized by Itskhoki (2008) and Antras, de Gortari and Itskhoki (2017) in an economy where entrepreneurs can decide whether to export or not. This makes labor supply decisions more elastic in an open economy, which may reduce redistribution at the optimum.

taxes should be imposed, what are the sufficient statistics that can guide the optimal taxation of new technology firms in practice? In second best environments—where income taxation is available, but taxes on specific factors are not—we have shown that there is a case for taxing new technology firms, if taxes on old technology firms are unavailable. We have derived multiple tax formulas as well as bounds on any Pareto efficient tax. We have also demonstrated how our tax formula could be combined with existing reduced-form evidence to compute the Pareto efficient tax on robots or trade.

Finally, we have used our formulas to conduct comparative statics. Through a simple example, we have illustrated that more robots and more trade, may go hand in hand with more inequality and lower taxes, in spite of robots and trade being responsible for the rise in inequality, and governments having extreme preferences for redistribution.

References

Acemoglu, Daron and Pascual Restrepo, "Robots and Jobs: Evidence from US Local Labor Markets," *mimeo MIT*, 2017.

Antras, Pol, Alonso de Gortari, and Oleg Itskhoki, "Globalization, Inequality and Welfare," *Journal of International Economics*, 2017, *108*, 387–412.

Arkolakis, Costas, Arnaud Costinot, and Andres Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

_ , _ , Dave Donaldson, and Andres Rodríguez-Clare, "The Elusive Pro-Competitive Effects of Trade," *Review of Economic Studies*, Forthcoming.

Autor, D., D. Dorn, and G. Hanson, "The China syndrom: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103*, 2121–2168.

Baeqee, David and Emmanuel Farhi, "Productivity and Misallocation in General Equilibrium," *mimeo Harvard University*, 2017.

Bagwell, Kyle and Robert W. Staiger, "An Economic Theory of GATT," *American Economic Review*, 1999, *89* (1), 215–248.

_ and _ , "Domestic Policies, National Sovereignty, and International Economic Institutions," *Quarterly Journal of Economics*, 2001, *116* (2), 519–562.

_ and _ , "The Economics of Trade Agreements in the Linear Cournot Delocation Model," *Journal of International Economics*, 2012, *88* (1), 32–46.

_ and _ , "Profit Shifting and Trade Agreements in Imperfectly Competitive Markets," *International Economic Review*, 2012, *53* (4), 1067–1104.

Basu, Susanto and John G. Fernald, "Aggregate Productivity and Aggregate Technology," *European Economic Review*, 2002, *46*, 963–991.

Bhagwati, Jagdish N., "Immesirizing Growth: A Geometrical Note," *Review of Economic Studies*, 1958, *25.*

_ , *The Generalized Theory of Distortions and Welfare*, Vol. Trade, Balance of Payments, and Growth: Papers in International Economics in Honor of Charles P. Kindleberger, Amsterdam: North-Holland, 1971.

Chetty, Raj, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 2012, *80* (3), 969–1018.

Chetverikov, Denis, Bradley Larsen, and Christopher John Palmer, "IV Quantile Regression for Group-Level Treatments, with an Application to the Effects of Trade on the Distribution of Wage," *Econometrica*, 2016, *84* (2), 809–833.

Diamond, Peter A., "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 1998, *88* (1), 83–95.

_ **and James A. Mirrlees**, "Optimal Taxation and Public Production II: Tax Rules," *American Economic Review*, 1971, *61* (3), 261–278.

Diamond, Peter and James Mirrlees, "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review*, 1971, *61* (1), 8–27.

Dixit, Avinash and Victor Norman, "Gains from Trade without Lump-Sum Compensation," *Journal of Intermational Economics*, 1986, *21*, 111–122.

Dixit, Avinash K. and Victor Norman, *Theory of International Trade*, Cambridge University Press, 1980.

Edgeworth, F.Y., "The theory of international values," *Economic Journal*, 1884, *4*, 35–50.

Feenstra, Robert C. and Tracy R. Lewis, "Trade adjustment assistance and Pareto gains from trade," *Journal of International Economics*, 1994, *36*, 201–222.

Galle, Simon, Andres Rodríguez-Clare, and Moises Yi, "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," 2017. Unpublished manuscript, UC Berkeley.

Graetz, Georg and Guy Michaels, "Robots at Work," *mimeo LSE*, 2018.

Grossman, Gene M. and Elhanan Helpman, "Protection for Sale," *American Economic Review*, 1994, *84* (4), 833–850.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles, "Should Robots Be Taxed?," *mimeo Northwestern University*, 2017.

Guesnerie, Roger, "Peut-on toujours redistribuer les gains a la specialisation et a l'echange? Un retour en pointille sur Ricardo et Heckscher-Ohlin," *Revue Economique*, 1998, *49* (3), 555–579.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura, "Income Taxation of U.S. Households: Facts and Parametric Estimates," *Review of Economic Dynamics*, 2014, *17* (4), 559–581.

Hosseini, Roozbeh and Ali Shourideh, "Inequality, Redistribution, and Optimal Trade Policy: A Public Finance," *mimeo Carnegie Mellon University*, 2018.

Hulten, Charles R., "Growth Accounting with Intermediate Inputs," *The Review of Economic Studies*, 1978, *45* (3), 511–518.

Itskhoki, Oleg, "Optimal Redistribution in an Open Economy," *mimeo Princeton University*, 2008.

Jacobs, Bas, "Optimal Inefficient Production," *mimeo Erasmus University*, 2015.

Mayer, Wolfgang and Raymond G. Riezman, "Endogenous Choice of Trade Policy Instruments," *Journal of International Economics*, 1987, *23*, 377–381.

— **and** — , "Tariff Formation in a Multidimensional Voting Model," *Economics and Politics*, 1989, *61-79.*

Mirrlees, J.A., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 1971, *38* (2), 175–208.

Naito, Hisairo, "Re-examination of uniform commodity taxes under non-linear income tax system and its implications for production efficiency," *Journal of Public Economics*, 1999, *1* (2), 65–88.

— , "Redistribution, production inefficiency and decentralized efficiency," *International Tax Public Finance*, 2006, *13*, 625–640.

Rodrik, Dani, *Political Economy of Trade Policy*, Vol. 3 of *Handbook of International Economics*, Elsevier, 1995.

Saez, Emmanuel, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68* (1), 205–229.

Scheuer, Florion and Ivan Werning, "The Taxation of Superstars," *Quarterly Journal of Economics*, 2017, *132*, 211–270.

Solow, Robert, "Technical Change and the Aggregate Production Function," *The Review of Economics and Statistics*, 1957, *39* (3), 312–320.

Spector, David, "Is it possible to redistribute the gains from trade using income taxation?," *Journal of International Economics*, 2001, *55*, 441–460.

Stiglitz, Joseph E., "Self-selection and Pareto efficient taxation," *Journal of Public Economics*, 1982, *17* (2), 213–240.

Tsyvinski, Aleh and Nicolas Werquin, "Generalized Compensation Principle," *mimeo Yale University*, 2018.

Waugh, Michael and Spencer Lyon, "Redistributing the Gains from Trade through Progressive Taxation," *mimeo NYU*, 2017.

Werning, Ivan, "Pareto Efficient Income Taxation," *mimeo MIT*, 2007.

A Proofs

A.1 Section (4.2)

Starting from equation (11), the equivalent variation, $EV(\phi_0, \phi_1)$, associated with a productivity shock from ϕ_0 to ϕ_1 can be computed as the unique solution to the differential equation

$$\frac{dEV(\phi,\phi_1)}{d\phi} = \frac{\partial G^*(y^*(\phi, EV(\phi,\phi_1));\phi)}{\partial \phi}, \text{ with initial condition } EV(\phi_1,\phi_1) = 0, \tag{27}$$

evaluated at $\phi = \phi_0$. Likewise, the compensating variation, $CV(\phi_0, \phi_1)$, can be computed as the unique solution to

$$\frac{dCV(\phi_0,\phi)}{d\phi} = \frac{\partial G^*(y^*(\phi, -CV(\phi_0,\phi));\phi)}{\partial \phi}, \text{ with initial condition } CV(\phi_0,\phi_0) = 0, \tag{28}$$

evaluated at $\phi = \phi_1$. In equations (27) and (28), $y^*(\phi, EV(\phi,\phi_1))$ and $y^*(\phi, -CV(\phi_0,\phi))$ are the counterparts to the compensated Hicksian demand functions in standard consumer theory, evaluated at the final and the initial utility level, respectively.

A.2 Section 5.3

The Lagrangian associated with the planner's problem (12) is given by

$$\mathcal{L} = \int U(\theta)d\Lambda(\theta) + \int \mu(\theta)\left(U'(\theta) + u_n(C(n(\theta),U(\theta)),n(\theta))n(\theta)\omega(\{p_i\},\{n(\theta)\})\right)d\theta$$
$$- \gamma G^*(\{c_j(\{q_i\},\{n(\theta)\},\{U(\theta)\}) - y_j(\{p_i\},\{n(\theta)\})\};\phi).$$

Integrating by parts, we get

$$\mathcal{L} = \int U(\theta)d\Lambda(\theta) - \int \mu'(\theta)U(\theta)d\theta + U(\bar{\theta})\mu(\bar{\theta}) - U(\underline{\theta})\mu(\underline{\theta})$$
$$+ \int \mu(\theta)u_n(C(n(\theta),U(\theta)),n(\theta))n(\theta)\omega(\{p_i\},\{n(\theta)\})d\theta$$
$$- \gamma G^*(\{c_j(\{q_i\},\{n(\theta)\},\{U(\theta)\}) - y_j(\{p_i\},\{n(\theta)\})\};\phi).$$

Since $U(\bar{\theta})$ and $U(\underline{\theta})$ are free we must have

$$\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0.$$

The first-order condition with respect to $U(\theta)$ leads to

$$\lambda(\theta) - \mu'(\theta) + \mu(\theta)u_{nC}(\theta)C_{U(\theta)}(\theta)n(\theta)\omega(\theta) - \gamma\nabla_{y^*}G^* \cdot c_{U(\theta)}f(\theta) = 0,$$

where λ denotes the density associated with Λ. From condition (7) and the normalization $p_1^* = 1$, we know that $p_i^* = G_{y_i}^*/G_{y_1}^*$ for all i. From the fact that $D_q c \cdot (p^* - q) = 0$ at the optimum, as argued in the main text, and the normalization $q_1 = 1$, we also know that $p_i^* = q_i$ for all i. Using these two observations, we get

$$\mu'(\theta) - \mu(\theta) u_{nC}(\theta) C_{U(\theta)}(\theta) n(\theta) \omega(\theta) = \lambda(\theta) - \gamma G_{y_1}^* \zeta(\theta) f(\theta),$$

with $\zeta(\theta) \equiv \sum p_j^* \frac{dc_j(\{q_i\}, C(n(\theta), U(\theta)))}{dU(\theta)}$. Let $\tilde{u}(C, x, \theta) \equiv u(C, x/w(\theta))$, $\tilde{u}_\theta \equiv \frac{\partial \tilde{u}}{\partial \theta}$, and $\tilde{u}_{\theta C} \equiv \frac{\partial^2 \tilde{u}}{\partial \theta \partial C}$. By definition, we have

$$\tilde{u}_\theta(\theta) = -u_n(\theta) n(\theta) \omega(\theta)$$

and

$$\tilde{u}_{\theta C}(\theta) = -u_{nC}(\theta) n(\theta) \omega(\theta)$$

Using the previous notation and using the fact that $C_{U(\theta)} = 1/u_C(\theta) = 1/\tilde{u}_C(\theta)$, we can rearrange the above first-order condition as

$$\mu'(\theta) \tilde{u}_C(\theta) + \mu(\theta) \tilde{u}_{\theta C}(\theta) = \tilde{u}_C(\theta)(\lambda(\theta) - \gamma G_{y_1}^* \zeta(\theta) f(\theta)). \tag{29}$$

Let $\tilde{\mu}(\theta) \equiv \mu(\theta) \tilde{u}_C(\theta)$. By definition, we also have

$$\tilde{\mu}'(\theta) = \mu'(\theta) \tilde{u}_C(\theta) + \mu(\theta) \tilde{u}'_C(\theta)$$

with

$$\tilde{u}'_C(\theta) = \tilde{u}_{\theta C}(\theta) + \tilde{u}_{CC}(\theta) C'(\theta) + \tilde{u}_{Cx}(\theta) x'(\theta),$$

which can be rearranged as

$$\tilde{u}'_C(\theta) = \tilde{u}_{\theta C}(\theta) + x'(\theta)[\tilde{u}_{CC}(\theta) \frac{dC}{dx} + \tilde{u}_{Cx}(\theta)]. \tag{30}$$

From agent θ's budget constraint, $e(q, C) = R(x)$, we know that

$$\frac{dC}{dx} = \frac{R'(x)}{e_C},$$

and from the first-order condition associated with (5), we know that

$$\frac{R'(x)}{e_C} = -\frac{\tilde{u}_x}{\tilde{u}_C}.$$

Combining the three previous equations, we obtain

$$\tilde{u}_C'(\theta) = \tilde{u}_{\theta C}(\theta) + \tilde{u}_C(\theta)x'(\theta)\rho(\theta), \tag{31}$$

where $\rho(\theta) \equiv \frac{\partial(\tilde{u}_x/\tilde{u}_C)}{\partial C} = \frac{\tilde{u}_{Cx}}{\tilde{u}_C} - \tilde{u}_{CC}\frac{\tilde{u}_x}{\tilde{u}_C^2}$ denotes the partial derivative, with respect to aggregate consumption, of the marginal rate of substitution between earnings and consumption. In turn, equations (29), (30), and (31) imply

$$\tilde{\mu}'(\theta) - \tilde{\mu}(\theta)\frac{dx}{d\theta}\rho(\theta) = \tilde{u}_C[\lambda(\theta) - \gamma G^*_{y_1}\zeta(\theta)f(\theta)].$$

Solving forward and using the fact that $\tilde{\mu}(\bar{\theta}) = 0$, we get

$$\begin{aligned}\tilde{\mu}(\theta) = &-\int_\theta^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)d\Lambda(z)\\ &+\gamma\int_\theta^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)G^*_{y_1}\zeta(z)dF(z).\end{aligned}$$

Since $\tilde{\mu}(\underline{\theta}) = 0$, we must also have

$$\gamma = \frac{\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)d\Lambda(z)}{\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)G^*_{y_1}\zeta(z)dF(z)},$$

which implies

$$\begin{aligned}\frac{\tilde{\mu}(\theta)}{\gamma G^*_{y_1}} = &\int_\theta^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)\zeta(z)dF(z)\\ &-\frac{\int_\theta^{\bar{\theta}} \exp[-\int_\theta^z \rho(v)dx(v)]\tilde{u}_C(z)d\Lambda(z)}{\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_{\underline{\theta}}^z \rho(v)dx(v)]\tilde{u}_C(z)d\Lambda(z)}\cdot\int_{\underline{\theta}}^{\bar{\theta}} \exp[-\int_{\underline{\theta}}^z \rho(v)dx(v)]\tilde{u}_C(z)\zeta(z)dF(z).\end{aligned}$$

Substituting into equation (13) and using the fact that $u_n(\theta)/u_C(\theta) = -(1-\tau(\theta))w(\theta)$, we obtain

$$D_p y\cdot(p^* - p) = \int \psi(\theta)(1-\tau(\theta))x(\theta)(\nabla_p\omega(\theta))d\theta,$$

with $\psi(\theta)$ given by equation (16), as argued in the main text.

A.3 Section 6.2

The first-order condition with respect to $n(\theta)$ is given by

$$\gamma\phi[y_{f,n(\theta)}(p_m,n) - h'(n(\theta)) + \frac{1}{\phi}y_{m,n(\theta)}(p_m,n\}]f(\theta) \tag{32}$$
$$= \mu(\theta)[h''(n(\theta))n(\theta) + h'(n(\theta))]\omega(p_m;\theta).$$

Since old technology firms choose their labor demand to maximize profits and agents choose their labor supply to maximize utility, we also know that

$$y_{f,n(\theta)}(p_m,n) + p_m y_{m,n(\theta)}(p_m,n) = w(\theta),$$
$$h'(n(\theta)) = w(\theta)(1-\tau(\theta)).$$

Thus, using the fact that $p_m^* = 1/\phi$, we can rearrange equation (32) into equation (22).

A.4 Section 6.4

Equation ((26)). Under the assumption that upper-level preferences are quasilinear, we have already argued that

$$\frac{\mu(\theta)}{\gamma\phi} = \Lambda(\theta) - F(\theta).$$

Together with equation (22), we therefore get

$$[w(\theta)\tau(\theta) + (p_m^* - p_m)y_{m,n(\theta)}] = \frac{[\Lambda(\theta) - F(\theta)]}{f(\theta)}h'(n(\theta))[\frac{\epsilon(\theta)+1}{\epsilon(\theta)}]\omega(p_m;\theta).$$

Using again the fact that $d\ln y_m(p_m, n(\theta);\theta)/d\ln n(\theta) = 1$ and $h'(n(\theta)) = w(\theta)(1-\tau(\theta))$, from the first-order condition of the agent's utility maximization problem, this leads to

$$\tau(\theta) = \tau^*(\theta) - \frac{p_m - p_m^*}{p_m}\frac{p_m y_m(\theta)}{x(\theta)}(1-\tau^*(\theta)), \tag{33}$$

with

$$\tau^*(\theta) \equiv \frac{1}{1 + \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{f(\theta)}{(\Lambda(\theta)-F(\theta))\omega(p_m;\theta)}}.$$

Combining the previous expression with Proposition 4, we obtain

$$\frac{t_m^*}{1+t_m^*} = \frac{\int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d\ln\omega(\theta)}{d\ln|y_m|} \cdot \tau^*(\theta) \cdot x(\theta)dF(\theta)}{p_m|y_m|[1 - \int \frac{\epsilon(\theta)}{\epsilon(\theta)+1} \cdot \frac{d\ln\omega(\theta)}{d\ln|y_m|} \cdot \tau^*(\theta) \cdot \frac{y_m(\theta)}{|y_m|}dF(\theta)]} \tag{34}$$

In the parametric example of Section 6.4, we have assumed

$$\epsilon(\theta) = \epsilon \text{ for all } \theta, \tag{35}$$

$$\Lambda(\theta) = 1 \text{ for all } \theta, \tag{36}$$

$$f(\theta) = 1 \text{ for all } \theta, \tag{37}$$

$$F(\theta) = \theta \text{ for all } \theta. \tag{38}$$

We therefore immediately get

$$\tau^*(\theta, p_m) = \frac{1}{1 + \frac{\epsilon}{\epsilon+1} \cdot \frac{1}{(1-\theta)\omega(p_m;\theta)}}. \tag{39}$$

In Section 6.4, we have also established that

$$w(p_m;\theta) = (1-\theta)^{-1/\gamma(p_m)},$$

which implies

$$\omega(p_m;\theta) = \frac{1}{\gamma(p_m)} \cdot \frac{1}{1-\theta}.$$

Substituting into equation (39), we therefore get

$$\tau^*(\theta) = \frac{1}{1 + \frac{\epsilon}{\epsilon+1}\gamma(p_m)} \equiv \tau^*. \tag{40}$$

In Section 6.4, we have also established that

$$\frac{d \ln \omega(p_m;\theta)}{d \ln p_m} = -\beta\gamma(p_m),$$

which implies

$$\frac{d \ln \omega(\theta)}{d \ln |y_m|} = -\beta\gamma(p_m)\frac{d \ln p_m}{d \ln |y_m(p_m, n)|} \equiv \frac{d \ln \omega}{d \ln |y_m|}. \tag{41}$$

Combining equations (34), ((35)), (40) and (41), we obtain

$$\frac{t_m^*}{1+t_m^*} = \frac{\frac{\epsilon}{\epsilon+1}\frac{d \ln \omega}{d \ln |y_m|}\tau^*}{1 - \frac{\epsilon}{\epsilon+1}\frac{d \ln \omega}{d \ln |y_m|}\tau^*(\theta, p_m)} \frac{\int x(\theta) dF(\theta)}{p_m |y_m|}, \tag{42}$$

Letting $s_m \equiv \frac{p_m|y_m|}{\int x(\theta)dF(\theta) + p_m|y_m|}$, this leads to equation ((26)).

Proposition 5. From equation ((26)), we know that

$$\frac{t_m^*}{1+t_m^*} = \frac{\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}\tau^*}{1-\frac{\epsilon}{\epsilon+1}\frac{d\ln\omega}{d\ln|y_m|}\tau^*}\frac{1-s_m}{s_m}$$

with

$$\frac{d\ln\omega}{d\ln|y_m|} = -\beta\gamma(p_m)\frac{d\ln p_m}{d\ln|y_m(p_m,n)|},$$
$$\tau^* = \frac{1}{1+\frac{\epsilon}{\epsilon+1}\gamma(p_m)},$$
$$s_m = \frac{p_m|y_m|}{\int x(\theta)dF(\theta)+p_m|y_m|},$$

This expression can be rearranged as

$$\frac{t_m^*}{1+t_m^*} = \frac{\Phi}{\rho-\Phi}\frac{1-s_m}{s_m} \tag{43}$$

with

$$\Phi = -\frac{\epsilon\beta\gamma(p_m)}{(\epsilon+1)+\epsilon\gamma(p_m)}, \tag{44}$$
$$\rho = \frac{\partial\ln|y_m(p_m,n)|}{\partial\ln p_m}.$$

We first demonstrate that Φ, s_m, and ρ can be expressed as functions of t_m^* and ϕ.

Using the fact that $p_m = (1+t_m^*)/\phi$, we can immediately rearrange equation (44) as

$$\Phi = -\frac{\epsilon\beta\gamma((1+t_m^*)/\phi)}{(\epsilon+1)+\epsilon\gamma((1+t_m^*)/\phi)} \equiv \Phi(t_m^*,\phi). \tag{45}$$

To express s_m and ρ as a function of t_m^* and ϕ, we further need to solve for the optimal labor supply of each agent, $n(\theta)$, which itself depends on the marginal income tax rates, $\tau(\theta)$. Together with equations (35)-(38), equation (33) implies

$$\tau(\theta) = \frac{\epsilon+1-\frac{t_m^*}{1+t_m^*}\frac{p_m y_m(\theta)}{x(\theta)}\gamma(p_m)}{\epsilon+1+\epsilon\gamma(p_m)}.$$

From the first-order condition of the old technology firms, we know that

$$\frac{p_m y_m(\theta)}{x(\theta)} = -\beta(\theta)\ln(1-\theta), \tag{46}$$

which leads to

$$\tau(\theta) = \frac{\epsilon + 1 + \frac{t_m^*}{1+t_m^*}\beta\gamma(p_m)\ln(1-\theta)}{\epsilon + 1 + \epsilon\gamma(p_m)}. \tag{47}$$

The optimal labor supply is given by the agent's first-order condition

$$n(\theta) = ((1-\tau(\theta))w(\theta))^{\epsilon}. \tag{48}$$

Combining equations (47) and (48) with the fact that $w(p_m;\theta) = (1-\theta)^{-1/\gamma(p_m)}$, we get

$$n(\theta) = (\frac{\gamma(p_m)}{\epsilon + 1 + \epsilon\gamma(p_m)})^{\epsilon}(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}(1-\theta)^{-\epsilon/\gamma(p_m)},$$

and in turn,

$$\int w(\theta)n(\theta)d\theta = (\frac{\gamma(p_m)}{\epsilon + 1 + \epsilon\gamma(p_m)})^{\epsilon}\int(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}\theta^{-\frac{1+\epsilon}{\gamma(p_m)}}d\theta$$

Using equation (46), we further get

$$p_m y_m(p_m, n(\theta);\theta) = -\beta\ln(1-\theta)(\frac{\gamma(p_m)}{\epsilon + 1 + \epsilon\gamma(p_m)})^{\epsilon}(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma(p_m)}}, \tag{49}$$

and in turn,

$$p_m|y(p_m,n)| = -\beta(\frac{\gamma(p_m)}{\epsilon + 1 + \epsilon\gamma(p_m)})^{\epsilon}\int\ln(1-\theta)(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma(p_m)}}d\theta. \tag{50}$$

The aggregate share of robots is therefore given by

$$s_m = \frac{\int\beta\ln(1-\theta)(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta}{\int(\beta\ln(1-\theta)-1)(\epsilon - \frac{t_m^*}{1+t_m^*}\beta\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta} \equiv s_m(t_m^*, \phi), \tag{51}$$

where we have again used $p_m = (1+t_m^*)/\phi$. The elasticity ρ can be computed in a similar manner. From equation (46) and the fact that $w(p_m;\theta) = (1-\theta)^{-1/\gamma(p_m)}$, we get

$$p_m y(p_m, n(\theta);\theta) = -\beta\ln(1-\theta)n(\theta)(1-\theta)^{-1/\gamma(p_m)}.$$

Using the previous expression with the definition of $\rho \equiv \frac{\partial\ln|y_m(p_m,n)|}{\partial\ln p_m}$, we get

$$\rho = \int\frac{y(p_m,n(\theta);\theta)}{|y_m(p_m,n)|}\frac{d\ln w(p_m;\theta)}{d\ln p_m}d\theta - 1.$$

Combining the previous expressions with equations (49), (50), and using the fact that $\frac{d\ln w(p_m;\theta)}{d\ln p_m} =$

$-\beta \ln(1-\theta)$ and $p_m = (1+t_m^*)/\phi$, we get

$$\rho = \frac{\int (\beta \ln(1-\theta) - 1) \ln(1-\theta)(\epsilon - \frac{t_m^*}{1+t_m^*}\beta \ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}} d\theta}{\int \ln(1-\theta)(\epsilon - \frac{t_m^*}{1+t_m^*}\beta \ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}} d\theta} \equiv \rho(t_m^*, \phi). \tag{52}$$

At this point, we have established that the three statistics in equation (43) can be expressed as $\Phi(t_m^*, \phi)$, $\rho(t_m^*, \phi)$, and $s_m(t_m^*, \phi)$. We can therefore rearrange equation (43) as

$$H(t_m^*, \Phi(t_m^*, \phi), \rho(t_m^*, \phi), s_m(t_m^*, \phi)) = 0,$$

with

$$H(t_m^*, \Phi, \rho, s_r) \equiv \frac{\Phi}{\rho - \Phi} \cdot \frac{1 - s_m}{s_m} - \frac{t_m^*}{1 + t_m^*}.$$

By the Implicit Function Theorem, we have

$$\frac{dt_m^*}{d\phi} = -\frac{dH/d\phi}{dH/dt_m^*}. \tag{53}$$

Since the tax on robots is chosen to maximize welfare, the second derivative of the government's value function, expressed as a function of t_m^* only, must be negative. Noting that H corresponds to its first derivative—which is equal to zero at the optimal tax—we therefore obtain

$$dH/dt_m^* < 0. \tag{54}$$

Since $\gamma(\cdot)$ is a strictly increasing function, equation (45) implies

$$\frac{\partial \Phi(t_m^*, \phi)}{\partial \phi} > 0. \tag{55}$$

To establish the monotonicity of s_m and ρ with respect to ϕ, it is convenient to introduce the following function:

$$d(t_m^*, \phi, \zeta; \theta) = (\epsilon - \beta \frac{t_m^*}{1+t_m^*} \ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}} (\ln(1-\theta))^{-\zeta}.$$

By construction, d is log-supermodular in (ϕ, ζ, θ). Since log-supermodularity is preserved by integration, the following function,

$$D(\phi, \zeta) = \int d(t_m^*, \phi, \zeta; \theta) d\theta,$$

is also log-supermodular. It follows that

$$\frac{D(\phi,\zeta=0)}{D(\phi,\zeta=-1)} = \frac{\int(\epsilon-\beta\frac{t_m^*}{1+t_m^*}\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta}{\int(\ln(1-\theta))(\epsilon-\beta\frac{t_r^*}{1+t_r^*}\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta} \text{ is increasing in } \phi,$$

$$\frac{D(\phi,\zeta=-2)}{D(\phi,\zeta=-1)} = \frac{\int(\ln(1-\theta))^2(\epsilon-\beta\frac{t_m^*}{1+t_m^*}\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta}{\int(\ln(1-\theta))(\epsilon-\beta\frac{t_r^*}{1+t_m^*}\ln(1-\theta))^{\epsilon}(1-\theta)^{-\frac{1+\epsilon}{\gamma((1+t_m^*)/\phi)}}d\theta} \text{ is decreasing in } \phi.$$

Noting that

$$s_m = \frac{1}{1-\frac{1}{\beta}\frac{D(\phi,\zeta=0)}{D(\phi,\zeta=-1)}},$$
$$\rho = \beta\frac{D(\phi,\zeta=-2)}{D(\phi,\zeta=-1)} - 1,$$

we obtain that

$$\frac{\partial s_m(t_m^*,\phi)}{\partial\phi} > 0, \tag{56}$$
$$\frac{\partial \rho(t_m^*,\phi)}{\partial\phi} < 0. \tag{57}$$

Since $\frac{\partial H}{\partial \Phi} < 0$, $\frac{\partial H}{\partial s_m} < 0$, and $\frac{\partial H}{\partial \rho} > 0$, inequalities (55)-(57) imply

$$\frac{dH}{d\phi} = \frac{\partial H}{\partial \Phi}\frac{\partial \Phi}{\partial \phi} + \frac{\partial H}{\partial s_m}\frac{\partial s_m}{\partial \phi} + \frac{\partial H}{\partial \rho}\frac{\partial \rho}{\partial \phi} < 0.$$

Combining this observation with equation (53) and (54), we conclude that $dt_m^*/d\phi > 0$.