

Distorted Innovation: Does the Market Get the Direction of Technology Right?*

Daron Acemoglu[†]

April 10, 2023

Abstract

In the presence of markup differences, externalities and other social considerations, the equilibrium direction of innovation can be systematically distorted. This paper builds a simple model of endogenous technology, which generalizes existing comparative static results and characterizes potential distortions in the direction of innovation. I show that empirical findings across a number of different areas are consistent with this framework's predictions and I use data from several studies to estimate its key parameters. Combining these numbers with rough estimates of differential externalities and markups, I provide suggestive evidence that equilibrium distortions in the direction of technology can be substantial in the context of industrial automation, health care, and energy, and correcting these distortions could have sizable welfare benefits.

Keywords: directed technological change, economic growth, energy, health care, innovation, technology.

JEL Classification: O33, O14, O31, J23, J24, C65, L65.

*This paper is based on the *American Economic Association Distinguished Lecture* delivered in New Orleans, January 2023. I thank Ufuk Akcigit, David Autor, Rena Conti, Joe Doyle, Michael Greenstone, Simon Johnson, Chad Jones, Amy Finkelstein, Will Rafey, Dani Rodrik and John Van Reenan for very useful comments, discussions and references. I am grateful to Juanita Jaramillo, Shinnosuke Kikuchi, Fredric Kong, and Todd Lensman for fantastic research assistance. Special thanks are due to David Hemous, Ralf Martin, Jacob Moscona and John Van Reenan for sharing data and help with the empirical work reported in this paper. Last but not least, this work heavily draws on joint work with several co-authors. I am particularly grateful to Ufuk Akcigit, David Autor, Simon Johnson, Pascual Restrepo and Fabrizio Zilibotti for their enduring contributions to my knowledge and understanding on these topics. All remaining errors are of course my own.

[†]Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, MA 02142. Email: daron@mit.edu

1 Introduction

There is broad agreement that technological change has been a major engine of economic growth and prosperity over the last 250 years.¹ However, not all innovations are created equal and the *direction of technology* matters greatly as well.

Both antibiotics and dietary supplements have resulted from new innovations and have led to products that have been consumed by billions of people around the world. But most would agree that antibiotics constitute a bigger technological breakthrough and have been socially more beneficial.² More strikingly, the same advances by early twentieth-century chemists, especially Fritz Haber and Carl Bosch, paved the way to both synthetic agricultural fertilizers, which massively boosted crop yields, and the large-scale production of more powerful explosives, which led to the deaths of millions of soldiers and civilians (e.g., Hager, 2009). Few people would think that these two advances have similar social value. Additionally, different technologies often create gains and losses for different groups and may even influence other major social outcomes, including civic participation and democracy.

Economists have long recognized that the overall amount of research effort may be insufficient and, as a result, government support for innovation, for instance in the form of investments in the research infrastructure or R&D tax credits, is beneficial (see, e.g., Jones and Williams, 1998, Bloom, Shankerman and Van Reenen, 2013, Howell, 2017, and Azoulay, Graff Zivin, Li and Sampat, 2019, for evidence). Nevertheless, a common perspective is that the market is the best judge of how research efforts should be allocated: once basic support to research is provided, the government should have a limited role in influencing the direction of innovation. There are indeed myriad examples of failed government attempts at “picking winners” (see Pack and Saggi, 2006, and Hufbauer and Jung, 2021, for reviews). The British science writer Ridley (2020) argues that the cumulative, step-by-step process of innovation is inevitably hampered when governments try to influence its direction. The opposite point of view emphasizes the myriad distortions in the equilibrium innovation process.

In this paper, I take an intermediate position. I assume that the market (working through competition between corporations and scientists) is best placed to experiment with new methods and carry out innovations, but it is possible for systemic factors to distort the direction of technology. This distinguishes my approach from both older-style industrial policy, where the government is assumed to have some ability to judge which sectors are more promising, and from more recent arguments claiming that the government can be as good as the private sector in innovation (e.g., Mazzucato, 2015).³

¹See, for example, Mokyr (1992, 2011) and Koyama and Rubin (2022).

²US annual expenditures are about \$10 billion in the 2010s for antibiotics and above \$30 billion for dietary supplements (see <https://academic.oup.com/cid/article/66/2/185/4093915>) and <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3952619/>). This is despite the fact that there are no well-established studies documenting the effectiveness of dietary supplements (see <https://pubmed.ncbi.nlm.nih.gov/32601065/> and <https://www.nytimes.com/2016/11/15/well/eat/studies-show-little-benefit-in-supplements.html>).

³Existing evidence suggests that although government encouragement to invest in high-tech sectors can have major benefits (e.g., Gruber and Johnson, 2019; Moretti, Steinwender and Van Reenen, 2021), top-down research often generates extensive distortions. For example, Howell, Rathje, Van Reenen and Wong (2021) show that traditional Defense Department research contracts have become less effective over time, while those that provide more open-ended support for new areas create more successful innovations. Similarly, convincing evidence of productivity benefits from industrial policy comes from settings in which such policy supported broad sectors, such as heavy and chemical industries in South Korea and Finland (Lane, 2022, Mitrunen, 2019). Additionally, Branstetter, Li and Ren (2022) show that recent Chinese industrial policy has not been successful in increasing firm productivity, while Acemoglu, Yang and Zhou (2022) provide evidence that top-down Chinese academic incentives have led to significant distortions in the direction of research.

To develop these ideas, I extend the directed technological change framework in Acemoglu (1998, 2002), focusing on an economy in which the private sector spearheads innovation and can target either one or both of two alternative, imperfectly substitutable technologies. From a positive perspective, the framework links the direction of technology to relevant market sizes (supplies of factors of production working with these technologies and consumer demand), the price of other inputs into the production process (for example, natural resources used in different sectors), markups and regulations. The implications of the framework are broadly in line with a growing body of empirical work, especially from sectors such as energy, health care, agriculture, modern automation and traditional industrial technologies.

More importantly for my focus here, the framework highlights several factors that can lead to *systematic* misalignment between market incentives and social objectives. First, some technologies generate negative externalities, while alternative paradigms aimed at performing similar production tasks may avoid these negative effects or even create positive externalities.⁴ The leading current example of this phenomenon is in the energy and transport sectors, where fossil-fuel-based energy creates carbon emissions and environmental damages, while renewables avoid such emissions. When the market does not price these damages, equilibrium innovation will be excessively directed towards fossil-fuel technologies. I argue that similar issues arise in health care, where some technologies, for instance those targeting prevention, may have greater social benefits than those aimed at high-tech procedures for late-stage cures. I will also suggest that the direction of technology may be distorted towards automation and away from worker-complementary technologies, because labor market imperfections create a wedge between the social cost of labor and the equilibrium wage.

Second, different sectors often have different markups, and I show that equilibrium incentives will be excessively biased towards higher-markup sectors and technologies.⁵ Health care illustrates this phenomenon, for curative technologies appear to have higher markups than preventative ones.⁶

Third, a variety of social forces may favor one paradigm ahead of alternatives. For example, the research community may value certain types of breakthroughs more than others, because they are viewed as the more exciting research area or because they are more useful for building a scientific reputation. One possible illustration may be from modern digital technologies, where many researchers believe that the most important and coveted advances are those that enable algorithms to reach human parity in a range of tasks. This perspective then creates greater incentives to work on automation rather than other paradigms aimed at more human-complementary tools.⁷

Fourth, when different technologies create distinct distributional effects and society cares about inequality (either for direct or indirect reasons related to political economy), the market will fail to inter-

⁴These positive externalities may also be on future research, for example, with some areas creating more substantial knowledge gains upon which future innovations can build. Another example of negative externalities would be “defensive innovations” undertaken by incumbents in order to prevent rivals from increasing their market share.

⁵High markups encourage more innovation effort but also simultaneously reduce the utilization of a technology. This latter effect implies that expanding the production level of high-markup technologies is also socially valuable. Yet this force is dominated by the former effect, and thus the equilibrium involves excessive innovation effort devoted to high-markup technologies.

⁶One telling example is emphasized in Howard, Bach, Brendt and Conti (2015): the melanoma drug Yervoy, approved in 2011, was marketed at the price of \$120,000 for a four-dose treatment by the pharmaceutical company Bristol-Myers Squibb. It extends life by about four months. I provide direct evidence on these markup differences in Section 5.

⁷See Acemoglu and Restrepo (2020b), Acemoglu, Jordan and Weyl (2022), Brynjolfsson (2022) and Acemoglu and Johnson (2023) for the argument that the general incentives in the artificial intelligence (AI) community and the agenda spearheaded by Alan Turing are creating an excessive focus on using AI technologies for automation.

nalize these additional considerations. Finally, the direction of innovation may be distorted because of coordination failures. For example, there may be insufficient diversity in research investments, or firms and innovators may fail to coordinate on more productive alternative paradigms, as I discuss briefly below.

While in practice all five of these effects are likely important, I focus on the first two, for two related reasons. First, these two channels can be, in principle, quantified by measuring markups or social benefits/externalities and are thus better candidates for “systematic distortions” in the direction of innovation. In the last part of the paper, I make a preliminary attempt at this type of quantification. Second, because these are quantifiable distortions, correcting them does not require government agencies to have superior information or an ability to “pick winners”. By comparison, it is more difficult to objectively determine whether an untried paradigm would be more successful or if the research community’s enthusiasm for a specific topic is a “fad” leading to excessive concentration of innovative effort.

I reestimate the empirical models from three studies in order to obtain some of the key parameters of my framework in these three different settings. These studies are: Acemoglu and Restrepo (2022) for research directed at automation technologies; Acemoglu, Moscona, Sastry and Williams (2023) for medical research directed towards different types of diseases; and Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen (2016) for fossil-fuel-based and cleaner innovations in the automobile industry. In each case, I identify the elasticity of substitution between different technologies and the degree to which past advances in one field create a relative advantage for the same field in the future. These two parameters are critical for both the equilibrium response of the direction of innovation to factor supplies, prices and policies, and for the divergence of equilibrium allocations from socially-optimal choices. I combine these numbers with estimates of the social costs/benefits of different technologies and markups to assess how distorted equilibrium technology choices are and the welfare gains from redirecting innovation. In each case, I provide suggestive evidence that innovation distortions and their welfare effects are sizable.

Related Literature. This paper builds on and extends the literature on directed technological change. The first explicit discussion of this topic is in Hicks’s (1932) argument that a high price for a factor induces technological change targeted at economizing on that factor. The induced technology literature of the 1960s explored whether technological change would be Harrod-neutral (purely labor augmenting), as typically imposed in neoclassical growth models (e.g., Kennedy, 1964, Ahmed, 1965, Samuelson, 1965, Drandakis and Phelps, 1966). But this literature relied on ad hoc rules determining the direction of technology and the exact form of these rules had defining effects on their results.

An early empirical investigation of these issues was Habakkuk’s (1962) seminal study of American technology in the nineteenth century. Habakkuk argued that the direction and speed of American technology was shaped by a desire to economize on scarce skilled labor in the country. Allen (2009) similarly suggested that Britain was the first country to industrialize because British labor was more expensive than labor in other European economies and in China.

The more recent literature on the direction of technological change has developed models with explicit R&D decisions targeted at different sectors and monopolistic profits from new technologies, which shape the composition of innovation. The first two papers within this area, Acemoglu (1998) and Kiley (1999), investigated the reasons why recent industrial technologies have often been skill-biased and why this skill bias may have accelerated starting in the 1980s, concurrently with the large increase in the supply of

educated workers in the US and other industrialized nations.

Acemoglu and Zilibotti (2001) develop a model of directed technological change in a multi-country setup, whereby technology choices in advanced economies are directed to their own needs, creating a form of “inappropriate technology” from the viewpoint of less developed nations with different factor endowments. Broader cross-country implications of directed technological change are studied in Gancia and Zilibotti (2005). Acemoglu (2003a) and Thoenig and Verdier (2003) explore how trade openness impacts the endogenous skill bias of technology. Acemoglu (2003b) and Jones (2006) investigate why an endogenous direction of technology can lead to Harrod-neutral advances and thus balanced growth as in textbook neoclassical models. Acemoglu and Linn (2004) and Costinot, Donaldson, Kyle and Williams (2019) study how demographic changes that alter the future market sizes of different types of medical technologies impact the direction of innovation, while in Acemoglu (2010) I present a formalization of the Habakkuk-Allen hypothesis, where labor scarcity can be a spur to faster economic growth. Bovenberg and Smulders (1995), Goulder and Schneider (1999), Di Maria and Valente (2006), Grimaud and Rouge (2008), Acemoglu, Aghion, Bursztyrn and Hémous (2012), Rodrik (2014), Acemoglu, Akcigit and Kerr (2016) and Hémous (2018), among others, discuss the balance between clean and dirty technologies and possible corrective policies in the presence of environmental externalities. Acemoglu and Restrepo (2018, 2022) and Hémous and Olsen (2022) explore the endogenous choice between automation and other types of technologies.

The model I present here generalizes Acemoglu (2002) in two important directions. First, I extend the baseline framework to perform comparative statics with respect to input prices, markups and externalities. Second, to the best of my knowledge, I undertake the first general analysis of the efficiency of the direction of technology within this framework, though Acemoglu, et al. (2012) provides a characterization of optimal policies to restore efficiency in a model of endogenous technology and carbon emissions.

Also closely related to this paper are a series of works that point out why the equilibrium direction of technology may be inefficient. Acemoglu (2011) shows that equilibrium innovation is often insufficiently diverse, investing too much in one of two alternative technologies. This innovation pattern then leaves the economy vulnerable to shifts in underlying conditions or blockages in existing technological paradigms. Acemoglu, Alp, Akcigit, Bloom and Kerr (2018) propose a model in which the distribution of innovation between firms of different sizes and ages is distorted.

Akcigit, Hanley and Serrano-Velarde (2021) distinguish between fundamental and applied research, and argue that the former generates more knowledge spillovers. The paper provides empirical evidence and a quantitative evaluation of this source of inefficiency. Distortions in the direction of technology resulting from different knowledge spillovers are also explored in Dechezleprêtre, Martin and Mohnen (2013) and Martin and Verhoeven (2022), and are present in models of international technology diffusion as well (e.g., Grossman and Helpman, 1993, Coe and Helpman, 1994). Acemoglu, Akcigit and Kerr (2016) show that new ideas in some fields matter more for subsequent innovation than others. These differential knowledge spillovers are complementary to the distortions emphasized in this paper.

Another related literature focuses on the choice between different technological paradigms and the possibility of inefficient lock-in (e.g., Dosi, 1982, and Arthur, 1989, as well as recent work by Acemoglu and Lensman, 2023).

The rest of the paper is organized as follows. Section 2 presents the basic framework of directed tech-

nological change, first in a static and then in a dynamic setting and characterizes equilibrium innovation. Section 3 compares the equilibrium allocation and the types of technologies developed to those that are socially optimal. Section 4 reviews several studies that provide evidence on the effects of market size, prices, markups and policies on the direction of innovation. Section 5 focuses on a few of these studies to obtain estimates of the key parameters of the framework. It then combines these parameters with numbers on markups and externalities to present a first evaluation of the extent of technology distortions and welfare gains from correcting them. Section 6 contains concluding comments, while the online Appendix presents additional derivations, results and details left out of the main text.

2 A Simple Model of Directed Technology

In this section, I provide a simple, two-sector model of directed technology. For simplicity, I start with a static setting and then outline the dynamic version, which is similar to the setup in Acemoglu (2002).

2.1 Static Environment

The economy is static and inhabited by a representative household with preferences given by

$$U = \ln C + \ln E, \tag{1}$$

where C denotes consumption, while E is an externality term, specified below. The log functional form is adopted to maximize the similarity with the infinite-horizon version of this model. There are three different types of labor (two of them working in the two sectors plus scientists). As is standard, I assume that labor income from all these types of labor accrues to the representative household.

The unique final good is produced with the production function

$$Y = \left[\gamma_1 Y_1^{\frac{\varepsilon-1}{\varepsilon}} + \gamma_2 Y_2^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{2}$$

where Y_1 and Y_2 denote the output levels of the two intermediate products, which are themselves produced using two different types of technologies (clean vs. dirty, preventative vs. curative, worker-friendly vs. automation, etc.). Their production functions are given by

$$Y_j = X_j^\alpha R_j^{1-\alpha}, \tag{3}$$

for $j \in \{1, 2\}$, where X_j denotes a variable input aggregate and R_j is a resource input, with exogenous price q_j^R .

The variable input is in turn produced with the following production function

$$X_j = \left(\int_0^{N_j} x_j(\nu)^{1-\beta} d\nu \right) \tilde{L}_j^\beta, \tag{4}$$

where \tilde{L}_j is a specialized factor employed only in sector j . For example, \tilde{L}_2 could correspond to skilled (or college-educated) labor, while \tilde{L}_1 could be unskilled labor.⁸ In addition, the $x_j(\nu)$'s denote the quantities

⁸In some applications, these could be the same labor allocated to the two sectors. In that case, L_1 and L_2 would be

of the different machine varieties used in the production of intermediate good $j \in \{1, 2\}$ and $\beta \in (0, 1)$. In this formulation, $[0, N_j]$ denotes the range of machines used in the production of $j \in \{1, 2\}$ and captures how advanced the technology for intermediate good j is. Once invented, each machine can be produced at the fixed marginal cost $\psi > 0$ in terms of the final good.

In what follows, I assume that all labor is supplied inelastically:

$$\tilde{L}_j = L_j \text{ for } j \in \{1, 2\}.$$

I model the *innovation possibilities frontier*, which specifies how new machine varieties are invented, by assuming that new ideas (or new machine varieties) are created by scientists. Specifically, the technology for creating new machine varieties is assumed to take the following static form:

$$N_j = \tilde{\eta}_j \phi(S_j) S_j, \tag{5}$$

where $\tilde{\eta}_j > 0$, S_j is the number of scientists assigned to technology j , and $\phi(S_j) = S_j^{\frac{\delta}{1-\delta}}$ with $\delta \in [0, 1)$. When $\delta > 0$, the innovation possibilities frontier features increasing returns to scale at the sectoral level. In the dynamic model, these increasing returns will take the form of *path dependence*, meaning that past advances in the technology of a sector will make further advances in the same sector easier. Scientists take the behavior of other scientists, and thus the value of the ϕ function, as given. The case of $\delta = 0$ corresponds to the useful benchmark in which there are no increasing returns to scale and scientists have a constant productivity in each sector.

Scientists that innovate and create varieties of machines become the owners of the technology monopolists that sell those varieties. This means that when a scientist invents a new machine for sector $j \in \{1, 2\}$, she will be able to make a profit π_j , which I characterize below.

The total number of scientists is fixed, so market clearing for scientists yields:⁹

$$S_1 + S_2 = \bar{S}. \tag{6}$$

To determine the profit levels from new machines for the two sectors, I adopt a simple market structure where each sector is subject to a fringe of competitive firms that can imitate and produce every machine, but do so less efficiently. This forces a limit price in each sector, given by

$$q_j = (1 + \mu_j)\psi, \tag{7}$$

where $\mu_j \in (0, \frac{\beta}{1-\beta}]$. This formulation provides a tractable form in which markups are potentially different between the two sectors.

endogenous and satisfy a single market-clearing constraint, $L_1 + L_2 \leq \bar{L}$. Equilibrium would then require their earnings in the two sectors to be equalize, i.e., $w_2/w_1 = 1$.

⁹Formally, in Appendix A, I suppose that each scientist has mass $\varepsilon > 0$, and then consider the limit where $\varepsilon \rightarrow 0$. This only matters in ensuring that deviations are well defined in the presence of externalities.

Additionally, note that the assumption that scientists take the value of $\phi(S_j)$ as given does not matter for the results, because with a fixed supply of scientists and the iso-elastic form of the ϕ function, the allocation of scientists between the two sectors is the same even if scientists form consortia that internalize the positive externalities they create on other scientists working in the same field. Finally, it is straightforward to make the supply of scientists endogenous to the income that scientists derive from innovation, but I will not do so in this paper.

Finally, I assume that the externality term in (1) takes a simple form given by

$$E = e^{-\sum_{j \in \{1,2\}} \tilde{\tau}_j \ln N_j}, \quad (8)$$

where $\tilde{\tau}_j \geq 0$ represents a negative externality from technology $j \in \{1,2\}$ (or if there are positive externalities, then $\tilde{\tau}_j < 0$). The assumption that the negative externalities originate from the level of technology is adopted for simplicity. Because these externalities do not impact market prices and are ignored by scientists and firms, they will play no role in the equilibrium allocation, but will have a major impact on the efficiency of the equilibrium.¹⁰

An equilibrium in this environment is defined as an allocation in which both the final good sector and the two intermediate sectors minimize costs, technology monopolists maximize profits by setting the limit price given in (7), scientists maximize their income by choosing which sector to innovate in, and all markets clear. I am particularly interested in the equilibrium level of relative technology, denoted by n^{EQ} (where $n \equiv N_2/N_1$).

2.2 Static Equilibrium

Cost-minimizing demands for machines and resources can be computed from the maximization problem

$$\max_{\{x_j(\nu), L_j, R_j\}} p_j \left(\int_0^{N_j} x_j(\nu)^{1-\beta} d\nu \cdot L_j^\beta \right)^\alpha R_j^{1-\alpha} - \int_0^{N_j} q_j(\nu) x_j(\nu) d\nu - w_j L_j - q_j^R R_j, \quad (9)$$

where w_j is the price (wage) of factor $j \in \{1,2\}$. Combining (7) with the expressions for machine and resource demands (provided in Appendix A), we obtain technology monopolists' profits as:

$$\begin{aligned} \pi_j(\nu) &= \mu_j \psi x_j(\nu) \\ &= \mu_j \psi \left[\left(p_j \left(\frac{(1-\beta)\alpha}{(1+\mu_j)\psi} \right)^\alpha \left(\frac{1-\alpha}{q_j^R} \right)^{1-\alpha} \right)^{\frac{1}{\alpha\beta}} L_j \right] \equiv \pi_j, \end{aligned} \quad (10)$$

where the square-bracketed expression in the second line is $x_j(\nu) \equiv x_j$, and the last equality defines the equilibrium flow profits for the two sectors, π_j , which, as claimed above, is identical for all machines used in sector $j = 1, 2$.

Setting the final product as the numeraire, the cost-minimization condition for the final good sector implies

$$p_j = \gamma_j \left(\frac{Y_j}{Y} \right)^{-\frac{1}{\varepsilon}}. \quad (11)$$

Combining this expression with (3), (4), and the expressions for machine and resource demands in

¹⁰If, instead, the externalities were from the production or consumption levels of the intermediates, technology would have additional indirect effects working through changes in equilibrium prices. The simplification enables me to remove these indirect effects.

Appendix A, we obtain:

$$p \equiv \frac{p_2}{p_1} = \left(\frac{\gamma_2}{\gamma_1} \right)^{\frac{\alpha\beta\varepsilon}{\sigma}} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{\frac{\alpha(1-\beta)}{\sigma}} \left(\frac{q_2^R}{q_1^R} \right)^{\frac{1-\alpha}{\sigma}} \left(\frac{N_2}{N_1} \right)^{-\frac{\alpha\beta}{\sigma}} \left(\frac{L_2}{L_1} \right)^{-\frac{\alpha\beta}{\sigma}}, \quad (12)$$

where $\sigma \equiv \alpha\beta\varepsilon + 1 - \alpha\beta$ is the *derived* elasticity of substitution between the two types of labor. Intuitively, the relative price of a sector's product is decreasing in the technology and labor supply to the sector, since these tend to expand its output. In addition, higher resource prices and markups increase a sector's relative price.

The price levels are then obtained by combining this relative price equation with the ideal price condition, which uses the fact that the final good is the numeraire:

$$[\gamma_2^\varepsilon p_2^{1-\varepsilon} + \gamma_1^\varepsilon p_1^{1-\varepsilon}]^{\frac{1}{1-\varepsilon}} = 1. \quad (13)$$

Factor prices are equal to the value of the marginal product of the relevant factor: $w_j = \alpha\beta p_j Y_j / L_j$ for $j = 1, 2$. Using this equation and (12), the relative wage of the two types of labor can be derived as

$$\frac{w_2}{w_1} = \left(\frac{\gamma_2}{\gamma_1} \right)^{\frac{\varepsilon}{\sigma}} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{-\frac{(1-\beta)(\sigma-1)}{\beta\sigma}} \left(\frac{q_2^R}{q_1^R} \right)^{-\frac{(1-\alpha)(\sigma-1)}{\alpha\beta\sigma}} \left(\frac{N_2}{N_1} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{L_2}{L_1} \right)^{-\frac{1}{\sigma}}. \quad (14)$$

This expression confirms that σ is indeed the elasticity of substitution between the two types of labor. Equation (14) additionally shows that N_2/N_1 plays the role of relative factor-augmenting technological change. We can see that $(\sigma - 1)/\sigma$ also regulates the impact of relative technology N_2/N_1 , markups and resource prices, since the net effect of these economic quantities depends on whether they affect the production level of an intermediate good by more or less than its price.

In an interior equilibrium in which research is directed to both technologies, scientists should make the same profits from improving the technology for either sector. Recalling that the productivity of a scientist when she works on technology j is $\tilde{\eta}_j \phi(S_j) = \tilde{\eta}_j S_j^{\frac{\delta}{1-\delta}}$, an interior equilibrium must satisfy: $\tilde{\eta}_1 S_1^{\frac{\delta}{1-\delta}} \pi_1 = \tilde{\eta}_2 S_2^{\frac{\delta}{1-\delta}} \pi_2$. Inverting $\phi(S_j)$, we have $N_j = \tilde{\eta}_j S_j^{\frac{1}{1-\delta}}$, or $S_j = (N_j / \tilde{\eta}_j)^{1-\delta}$. Substituting this into the condition for an interior equilibrium and defining $\eta_j \equiv \tilde{\eta}_j^{1-\delta}$, we have

$$\eta_1 N_1^\delta \pi_1 = \eta_2 N_2^\delta \pi_2. \quad (15)$$

Combining this equation with (10) and (12), we obtain

$$n^{EQ} = \left[\frac{\eta_2}{\eta_1} \left(\frac{\gamma_2}{\gamma_1} \right)^{\frac{\varepsilon}{\sigma}} \frac{\mu_2}{\mu_1} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{-\frac{\sigma-(1-\beta)}{\beta\sigma}} \left(\frac{q_2^R}{q_1^R} \right)^{-\frac{(\sigma-1)(1-\alpha)}{\alpha\beta\sigma}} \left(\frac{L_2}{L_1} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\delta\sigma}}. \quad (16)$$

Equation (16) links the equilibrium technology ratio between the two sectors to parameters of the final good production function, the innovation possibilities frontier, resource prices, markups, and the relative supplies of factors employed in the two sectors. For example, focusing on the case where $\delta\sigma < 1$, n^{EQ} is increasing in L_2/L_1 if and only if $\sigma > 1$, as I discuss in greater detail below.

The next proposition follows from this discussion and equation (16), and the uniqueness of equilibrium

is established in Appendix A.

Proposition 1 *Suppose that $\delta < 1/\sigma$. Then there exists a unique equilibrium in which the relative technology ratio is given by (16).*

The comparative statics of the direction of technology in this unique equilibrium are provided readily by equation (16) and will be discussed after I present the dynamic version of this environment in the next subsection.

To understand the role of the condition $\delta < 1/\sigma$, note that the stabilizing economic force in this model is the lower price of the intermediate good that is technologically more advanced as shown by (12), and this force is stronger when σ is lower. The destabilizing force, on the other hand, is the extent of increasing returns, δ . When $\delta < 1/\sigma$, the sector that is further ahead technologically faces sufficiently lower returns from innovation and this ensures the existence and uniqueness of the interior equilibrium.

In contrast, when $\delta > 1/\sigma$, the degree of increasing returns to scale in research is sufficiently strong that there does not exist an equilibrium in which research is directed towards both sectors. I show in Appendix A that in this case there are two corner equilibria—all scientists working in sector 1 or all scientists working in sector 2. This discussion also illustrates why the comparative statics of (16) are only relevant when $\delta < 1/\sigma$.

2.3 Dynamic Environment

I now present the dynamic version of this economy, for brevity emphasizing only the elements that are different from the static setup. Suppose that time is continuous and runs to infinity. There is an infinitely-lived representative household with preferences given by $U(0) = \int_0^\infty e^{-\rho t} U(t) dt$, with $U(t) = \ln C(t) + \ln E(t)$, where $C(t)$ denotes consumption at time t , $E(t)$ is the externality term as in the text, and ρ is the discount rate of the representative household. Analogously with (8), we have $E(t) = e^{-\sum_{j \in \{1,2\}} \bar{\tau}_j \ln N_j(t)}$.

All of the equilibrium conditions derived in the static model now apply, except that they should be indexed by time. The main difference is the innovation possibilities frontier, which takes the form

$$\dot{N}_j(t) = \eta_j N_j(t)^{(1+\delta)/2} N_{\sim j}(t)^{(1-\delta)/2} S_j(t), \quad (17)$$

where $S_j(t)$ denotes scientists working for innovation in technology j at time t and the $N_j(t)^{(1+\delta)/2}$ term captures path dependence in innovation from one's own sector, while the $N_{\sim j}(t)^{(1-\delta)/2}$ is the contribution of the technology of the other sector. This innovation possibilities frontier is the dynamic analogue of (5). Notice the difference from the increasing returns to scale in (5): in the dynamic case, there is a form of increasing returns to scale, but it is realized over time. This is the reason I refer to δ as the degree of *path dependence*—when $\delta > 0$, once a sector is technologically ahead of the other one, it becomes more productive in generating new innovations.

The total number of scientists is again fixed, so $S_1(t) + S_2(t) = \bar{S}$ at all t . Scientists that innovate and create new varieties now become the perpetual owners of the technology monopolists that sell those varieties. Suppose also that resource prices, the q_j^R 's, are constant, which implies that profits from technology $j = 1, 2$ in this dynamic environment are constant and are still given by π_j in equation (10).

I first focus on an interior balanced growth path (BGP) in which $n(t) \equiv N_2(t)/N_1(t)$ is constant and thus scientists work on both technologies. This requires

$$\eta_1 \pi_1 N_1(t)^\delta = \eta_2 \pi_2 N_2(t)^\delta \text{ for all } t.$$

This condition is identical to (15) in the previous subsection. Hence the BGP technology ratio in this dynamic model is identical to the equilibrium technology ratio in the static model.

Proposition 2 *There exists a unique BGP, where the equilibrium direction of technology is given by equation (16).*

The fact that the unique BGP ratio coincides with the static equilibrium technology ratio is because of the way in which the static model was set up to mimic the insights of the dynamic framework.¹¹ While the BGP here coincides with the equilibrium of the static model, the full equilibrium path of the dynamic model leads to somewhat different results, as explained in the next proposition.

Proposition 3 *If $\delta < 1/\sigma$, the unique interior BGP (given by (16)) is globally (saddle-path) stable. In particular, starting from any initial conditions, the economy tends to this interior BGP. Moreover, the unique dynamic equilibrium allocates all scientists to the sector that is relatively behind (compared to the BGP). As a result, the BGP is reached in finite time.*

If $\delta > 1/\sigma$, then the interior BGP is unstable, and starting from almost all initial conditions the economy limits to an allocation in which only one of the two technologies advances.

This proposition clarifies why the case with $\delta < 1/\sigma$ is the focal one in the dynamic economy as well, and the intuition for this condition is similar: the stabilizing force via relative price changes should be stronger than the destabilizing force due to path dependence.

When $\delta > 1/\sigma$, the equilibrium in the dynamic model is still unique (in contrast to the static model, where there were multiple equilibria), but now the relative technology level identified by condition (16) corresponds to an unstable BGP, and the economy will never converge to it. Rather, the equilibrium allocation will limit to one of the two corner BGPs, where the economy has a constant growth rate, driven by research in only one of the two technologies.

2.4 Some Properties of Equilibrium Technology Choices

I now review some properties of equilibrium direction of technology (using either the static equilibrium or the BGP of the dynamic equilibrium). This discussion will be brief because most of this material is familiar from previous work and is not my main focus here, though recognizing these comparative statics helps build intuition about the workings of the model.

Relative supply effects: The direction of technology is determined by the relative supply of labor used with the two types of technologies, L_1 and L_2 . As in Acemoglu (1998, 2002) the implications of relative supplies on the direction of technology depend on market size and price effects. Holding prices

¹¹The existence of a unique BGP is a consequence of the simplifying functional form assumptions. In general, as discussed in Acemoglu (2007), multiple equilibria are possible. But given my focus here, uniqueness enables me to focus on issues of distorted technology more directly.

(in particular p_1 and p_2) fixed, greater relative supply of one type of labor expands the market size of the technology complementing that type of labor and it further encourages the development of this complementary technology (i.e., a higher L_2/L_1 increases n^{EQ}). However, in equilibrium, prices also adjust and this creates a countervailing force. Whether this countervailing force is more powerful than the direct market size effect depends on the elasticity of substitution between the two types of labor. Specifically, when $\sigma > 1$, the market size effect dominates the price effect, and n^{EQ} is increasing in L_2/L_1 . In contrast, when $\sigma < 1$, the price effect is more powerful and n^{EQ} is decreasing in L_2/L_1 .

Weak bias of technology: As we have just seen, the impact of L_2/L_1 on n^{EQ} is ambiguous and depends on the elasticity of substitution between the two factors. Nevertheless, as emphasized in Acemoglu (2002), there is a general, unambiguous result about the bias of technology. A change in technology is said to be *biased towards a factor* if, holding all other variables constant, it increases the relative price of that factor. The main result that holds in this class of models is that an increase in L_2/L_1 always induces a change in technology that is (weakly) biased towards L_2 . For example, if the relative supply of college-educated workers increases, then technology becomes more skill-biased. Intuitively, this is because, as equation (14) demonstrates, when $\sigma > 1$, a greater L_2/L_1 raises n^{EQ} , and in this case it is also a higher level of n^{EQ} that is biased towards type 2 workers. Conversely, when $\sigma < 1$, it is a *decrease* in n^{EQ} that is biased towards type 2 workers, and in this instance, higher L_2/L_1 leads to *lower* n^{EQ} . Hence, regardless of the exact value of the elasticity of substitution between the two factors, technology always (weakly) moves in a direction that is favorable to the more abundant factor. Among other things, this force might explain why aggregate technology has become more skill-biased over the last eight decades, while the supply of skilled workers in the industrialized world has risen rapidly (Acemoglu, 1998). Acemoglu (2007) shows that this weak bias result is more general and holds without any of the functional form assumptions imposed here, provided that some mild regularity conditions are satisfied.

Strong bias of technology: By substituting the expression for n^{EQ} from (16) into (14), we obtain the long-run (endogenous technology) relationship between relative supplies and relative wages as

$$\left(\frac{w_2}{w_1}\right)^{BGP} = \Gamma \left(\frac{L_2}{L_1}\right)^{\frac{\sigma-2+\delta}{1-\delta\sigma}}, \quad (18)$$

with

$$\Gamma \equiv \left[\frac{\eta_2}{\eta_1} \left(\frac{\gamma_2}{\gamma_1}\right)^{\frac{\varepsilon(1-\delta)}{\sigma-1}} \left(\frac{\mu_2}{\mu_1}\right) \left(\frac{1+\mu_2}{1+\mu_1}\right)^{-\frac{1-\delta(1-\beta)}{\beta}} \left(\frac{q_2^R}{q_1^R}\right)^{-\frac{(1-\alpha)(1-\delta)}{\alpha\beta}} \right]^{\frac{\sigma-1}{1-\delta\sigma}},$$

where recall that $\delta\sigma < 1$. This equation implies that the relationship between relative wages and relative supplies is upward-sloping when $\sigma > 2 - \delta$, exactly as in Acemoglu (1998, 2002). Intuitively, the condition $\sigma > 2 - \delta$ ensures that technology moves sufficiently in the direction of the factor that becomes more abundant. With this powerful change in the direction of technology, the demand for the more abundant factor increases so much that the overall consequence is to raise this factor's marginal product more than that of the less abundant factor. Consequently, the locus of long-run equilibria becomes upward sloping—greater relative supply translates into greater relative wage. Notice that, when technology is fixed, relative demand curves are always downward sloping in this model (as in all models with price-taking firms). The upward-sloping demand curve is a consequence of technology's response to changes

in relative supplies of factors. Acemoglu (2007) provides a version of the same result for more general technologies and also an analogue of this result for the wage level of a factor rather than its relative wage.

Resource prices: Equation (16) also clarifies that resource prices will have a major impact on the direction of technology. In particular, when $\sigma > 1$, an increase in q_2^R (relative to q_1^R) reduces n^{EQ} , because higher resource prices for sector 2 make production and thus the technologies being used in this sector less profitable.

Effects of markups: Equation (16) further highlights a first-order effect of markups on the direction of technology, and under relatively weak conditions, a higher μ_2 (holding μ_1 constant) increases n^{EQ} .¹²

Many of these theoretical implications receive empirical support, as I discuss in Section 4.

3 Distorted Technology

This section compares the equilibrium and socially optimal technology choices and identifies several reasons why equilibrium technology choices will be distorted.

3.1 Socially Optimal Direction of Technology

I now consider the social planner’s solution in the static environment (the same exercise for the dynamic setup is presented in Appendix A). Differently from equilibrium incentives, the social planner takes into account the externalities that the two intermediates generate. Naturally, the planner also cares about the full income stream accruing to the representative household, rather than just the monopoly profits.

In what follows I further focus on the case in which the social planner cannot directly control prices and allocations—and thus will not be able to correct for externalities and markups by introducing Pigovian taxes/subsidies. This choice has three motivations. First, practical (information-related) or political constraints often prevent governments from removing monopoly markups or may even make it difficult to implement corrective taxes for externalities. Second, as discussed in Acemoglu, et al. (2012), Pigovian taxes are not always sufficient by themselves to restore optimality when the direction of technology is endogenous.¹³ Third, this choice also enables me to clearly focus on the distortions created by the allocation of research effort and the welfare gains from eliminating these technology distortions (rather than the full welfare consequences of various microeconomic distortions).

Given these assumptions, the only choice of the social planner is the allocation of scientists between the two technologies. In practice, this can be achieved by targeted research subsidies or regulations, and here I assume that the planner directly controls this allocation. Hence, in the static environment, the planner’s problem can be written as maximizing (1) by choosing S_1 and S_2 subject to (6) and the innovation possibilities frontier (5), and taking all other equilibrium relationships, and in particular the

¹²As discussed in the Introduction, the countervailing force here is that higher markups reduce output and via this channel increase prices. It is straightforward to verify that more research is directed to sector $j = 1, 2$ when its markup μ_j increases, provided that $\sigma\beta + \mu_j(1 - \sigma)(1 - \beta) > 0$. This condition is satisfied whenever $\sigma \leq 1$ or whenever μ_j is not too large.

¹³This is because in models with endogenous innovation, there are distortions both in the production sector (captured by the externalities targeted by Pigovian taxes) and in the allocation of research effort between different sectors (due to monopoly profits and knowledge externalities, such as the path dependence introduced above). As a result, optimal allocations should correct for both sets of distortions. For example, in the context of the energy sector, relying just on carbon taxes without actively redirecting technological change away from fossil-fuels would slow down the transition to clean energy and amplify its short-run costs.

price function (12), as given. This yields a simple maximization problem for the social planner:

$$\max_{S_1, S_2 \geq 0: S_1 + S_2 \leq \bar{S}} \ln Y[N_1, N_2] + \ln E[N_1, N_2]$$

subject to (5), (6) and (12). Taking the first-order conditions for this expression, noting that $d \ln N_j = dN_j/N_j$, and substituting for S_j in terms of N_j as in the equilibrium analysis, this necessary condition for an interior social optimum can be written as

$$\eta_1 \left[\frac{d \ln Y}{d \ln N_1} + \frac{d \ln E}{d \ln N_1} \right] = n^{-(1-\delta)} \eta_2 \left[\frac{d \ln Y}{d \ln N_2} + \frac{d \ln E}{d \ln N_2} \right]. \quad (19)$$

Clearly, $d \ln E/d \ln N_j = -\tilde{\tau}_j$, and in Appendix A I prove that $d \ln Y/d \ln N_j = \gamma_j^\varepsilon p_j^{1-\varepsilon}$. Moreover, defining $\tau_j \equiv \tilde{\tau}_j/(\gamma_j^\varepsilon p_j^{1-\varepsilon})$ as a price-adjusted externality, the first-order condition can be simplified to

$$\eta_1 \gamma_1^\varepsilon p_1^{1-\varepsilon} (1 - \tau_1) = \eta_2 p_2^{1-\varepsilon} \gamma_2^\varepsilon (1 - \tau_2) n^{-(1-\delta)}.$$

We can then substitute from (12), and solve for the socially optimal ratio of technology between the two sectors, n^{SP} , as

$$n^{SP} = \left[\frac{\eta_2}{\eta_1} \left(\frac{\gamma_2}{\gamma_1} \right)^{\frac{\varepsilon}{\sigma}} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{\frac{1-\beta}{\beta} \frac{1-\sigma}{\sigma}} \left(\frac{1 - \tau_2}{1 - \tau_1} \right) \left(\frac{q_2^R}{q_1^R} \right)^{-\frac{(\sigma-1)(1-\alpha)}{\alpha\beta\sigma}} \left(\frac{L_2}{L_1} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\delta\sigma}}. \quad (20)$$

It is also useful to write the ratio of socially optimal and equilibrium technologies as

$$\frac{n^{SP}}{n^{EQ}} = \left[\left(\frac{\mu_2}{\mu_1} \right)^{-1} \left(\frac{1 + \mu_2}{1 + \mu_1} \right) \left(\frac{1 - \tau_2}{1 - \tau_1} \right) \right]^{\frac{\sigma}{1-\delta\sigma}}. \quad (21)$$

It can be verified that, given τ_1 and τ_2 , a higher μ_2 always implies a lower n^{SP}/n^{EQ} .¹⁴ There are indirect effects from markups, but the overall impact from a higher sectoral markup is to distort technology towards that sector. Additionally, a higher τ_2 always implies a lower n^{SP} and n^{SP}/n^{EQ} , because of the negative externalities. Finally, the impact of all of these factors on the extent of technology distortion is amplified by $\sigma/(1 - \delta\sigma)$. This is because a higher elasticity of substitution between factors and a greater degree of increasing returns to scale (or path dependence) in innovation makes the equilibrium direction of technology more responsive to markups and the social planner's preferred direction more sensitive to externalities. The next proposition summarizes these results.

Proposition 4 *Suppose that $\delta < 1/\sigma$. Then the social planner's problem has a unique solution given by (20). Greater externalities and higher markups in sector j imply that equilibrium technology is excessively distorted towards sector j .*

This proposition implies that the only sources of divergence between the equilibrium and the social

¹⁴Recall that τ_1 and τ_2 are functions of prices, and in (20), they are evaluated at the relative technology level n^{SP} . However, the interpretation of (21) requires some caution. In particular, in writing this expression we have to hold τ_1 and τ_2 fixed. Or alternatively, when distortions are small, intermediate prices p_1 and p_2 under n^{EQ} and n^{SP} will be approximately the same and thus a given level of $\tilde{\tau}_j$ will map to approximately the same level of τ_j .

planner’s solution in this setting are due to markups and externalities. The social planner would like to move equilibrium technology away from sectors that have high markups and high negative externalities.

Equation (21) additionally implies that technology distortions can be quantified using four sets of quantities: markup differences, μ_2/μ_1 ; externality differences, $(1 - \tau_2)/(1 - \tau_1)$; the degree of increasing returns to scale, δ ; and the elasticity of substitution between the two types of labor used in the two sectors, σ (which is in turn a function of ε , α and β).

As in the equilibrium characterization in Proposition 1, Proposition 4 focuses on the case where $\delta < 1/\sigma$. When this condition is violated, the social planner prefers all scientists to work only in one of the two sectors (see Appendix A), whereas, as we have seen, all scientists working in either sector is an equilibrium.

Given the equilibrium characterization, one can also compute the welfare loss in equilibrium relative to the social optimum. Appendix A provides a first-order approximation to the change in welfare between the equilibrium and the social optimum, and I present estimates of this welfare loss in the context of the applications in Section 5.

3.2 Other Considerations

Before moving to an assessment of the quantitative extent of distorted technology in various applications, I comment on a few additional issues.

First, I simplified the discussion by ignoring other sources of distortions in the direction of technology. One potentially important type of distortion originates from visions, beliefs, fads and ideologies. For example, the private sector may come to believe that only one path of development of a scientific platform is feasible, or may be gripped by a “technology fad”. These issues are discussed in Acemoglu and Restrepo (2020b) and Acemoglu and Johnson (2023) in the context of artificial intelligence (AI)—arguing that the influence of dominant companies and certain research approaches developed in the 1950s and 60s pushed the field too much towards automation-related applications of AI. These considerations can be introduced in the current model in a reduced-form manner by assuming that the market’s assessment of η_1 and η_2 are systematically biased away from the true values of these parameters. Alternatively, one of the sectors may offer greater reputation-building opportunities to researchers. The more interesting question, which is beyond the scope of the current paper, is how such misperceptions or distorted incentives arise and whether there could be systematic ways in which government regulation could detect and prevent them.

Second, for tractability’s sake, I have assumed that the degree of increasing returns to scale, captured by the parameter δ , is the same in the two sectors. In practice, certain types of research, for example those targeting a scientific breakthrough, or the “research” rather than the “development” part of R&D, may generate more knowledge spillovers (e.g., Akcigit et al., 2021). Such spillovers can also be introduced in our context, though measuring the exact extent of such externalities is challenging.

Third, policymakers may also wish to take into account distributional and other social effects.¹⁵ If society engages in costly fiscal redistribution in order to increase the incomes of certain groups (e.g., the unemployed, low-skill workers, etc.), then we can think of technologies that directly increase these

¹⁵Inequality generated by some technologies may create additional social problems (as argued, for example, by Wilson, 1996, and documented by Autor, Dorn and Hanson, 2021), or may erode support for democracy (as shown in Acemoglu, Ajzeman, Aksoy, Fiszbein and Molina, 2021). These considerations would constitute additional reasons for altering the direction of technology. Since these effects are harder to quantify, they fall beyond the scope of the current paper.

groups' productivity as generating first-order pecuniary externalities, which can again be captured by our τ parameters.

Fourth, there may be reasons why the market underinvests in diverse technologies, as argued in Acemoglu (2011). Specifically, when there are shifts in which technologies are appropriate in different time periods, the market economy may underinvest in having a diverse portfolio of technologies that can act as a stepping stone when the underlying environment changes.

Finally, in richer models, there can be coordination failures whereby the market coordinates on or stays too long with an inferior technology (see Acemoglu and Lensman, 2023, for recent work on this topic). Once again, quantifying distributional, diversity, and coordination effects is more challenging, and I leave these issues for future work as well.

4 Existing Empirical Evidence

In this section, I review several empirical papers from the area of energy, health technologies, agriculture, modern automation technologies, and the introduction of new industrial machinery during the Industrial Revolution to provide an overview of a body of growing empirical evidence on how market sizes, resource prices and policy impact the direction of technology. The available evidence generally supports the predictions of the theoretical framework of this paper.

4.1 Energy

There is a large and growing literature that shows the responsiveness of both energy-generation and energy-use technologies to resource prices. Newell, Jaffe and Stavins (1999) studied the impact of energy prices on energy-saving innovations. These authors collected data on the cooling/heating capacity, energy flow, energy efficiency and price of room air conditioners, central air-conditioning units, and gas water heaters from the Sears-Roebuck catalogs between 1957 and 1993. Their results show that higher energy prices have a significant impact on energy efficiency—the models offered to consumers became more energy-efficient when resource costs rose. The authors also present some evidence that energy standard regulations had a similar effect for room air conditioners. Consistent with the idea that there is a strong trade-off between different types of technologies, the authors additionally show that energy efficiency adjustments are associated with higher prices, and in fact, they do not find significant effects on the overall amount of technological change. Hence, this study suggests that the direction of technology may be more responsive to resource prices than the overall amount of technological change, which is consistent with the framework presented here when $\sigma > 1$.

Popp (2002) studies US patent and citations data from 1970 to 1994. He establishes a robust association between energy prices and energy-efficient innovations. Popp also shows a significant role of the knowledge base, reminiscent of path dependence in the innovation possibilities frontier above.

Aghion et al. (2016) provide additional evidence consistent with these patterns. These authors build a firm-level data set of automobile-related patents across 80 countries, and classify these innovations into dirty and clean technologies—e.g., internal combustion engine vs. hybrid and electric vehicles. They show that higher fuel prices induced by carbon taxes lead to more clean and less dirty innovations in the automobile industry. They also estimate statistically significant path dependence. In Section 5, I use this

study’s data to present some related results as a basis of my quantitative exercise.

More recent work by Acemoglu, Aghion, Barrage and Hémous (2022) documents a relationship between natural gas prices, driven by the US shale gas boom, and overall green patenting (relative to either all patents, energy patents or dirty patents). In particular, green patents surged when natural gas prices were high and then declined as the shale gas boom kicked in.

Overall, the evidence from the energy sector is fairly clear that resource prices have the expected impact on the direction of technology—and the direction of technology is possibly more responsive than, the overall amount of innovation. There is also evidence of path dependence, whereby energy-efficient (or green) innovations build on a specific knowledge base that past innovations of this type have created.

4.2 Health and Medical Technologies

The direction of health care and medical technologies appears to be highly responsive to market sizes, prices and regulations, along the lines of the predictions of the framework presented here. Finkelstein’s pioneering (2004) study focuses on several policy changes, expanding the market size for certain vaccines. Specifically, in 1991 the Center for Disease Control recommended that all infants be vaccinated against Hepatitis B, while in 1993 Medicare began covering the full cost of influenza vaccination for Medicare recipients (without any copayments). Finkelstein also looks at a 1986 reform indemnifying manufacturers from lawsuits from potential adverse reactions to childhood vaccines against polio, diphtheria-tetanus, measles mumps and rubella (MMR) and pertussis. Finkelstein estimates a 2.5-fold increase in the likelihood of clinical trials for the relevant vaccines following the policy-induced expansion of market size.¹⁶

Acemoglu and Linn (2004) focus more directly on the market size for new pharmaceuticals. They exploit variation originating from demographic change—for example, the baby boomer generation first creating demand for pharmaceuticals targeted at younger and middle-aged patients, and later as this cohort aged, for drugs targeting diseases for older patients. They find a powerful impact of market size on the introduction of new molecular entities, as well as the entry of new generics. Their baseline estimate suggests that a 1% increase in market size is associated with a 4% increase in new non-generic drugs. In subsequent work Acemoglu, Cutler, Finkelstein and Linn (2006) provide suggestive evidence that Medicare induced an increase in pharmaceutical innovations targeted at the elderly. Costinot, Donaldson, Kyle and Williams (2019) provide similar evidence from a cross-country setting. These authors combine predictions about the direction of innovation with the home market effect (whereby countries specialize in and export products targeted at their home market), and document that countries invest more in and export drugs that have a greater demand among their home population.

More recent research by Acemoglu, et al. (2023) assembles a comprehensive data set of cross-country medical research and disease burdens impacting different countries. They estimate a strong association between the burden from a disease and research directed towards that disease. Below I also present regression evidence from the data set compiled by these authors.

Finally, the only paper I am aware of that provides evidence relevant for the effects of markup differences is Budish, Roin and Williams (2015). These authors observe that the US patent system,

¹⁶Finkelstein does not find an increase in medical trials and patents, which may be due to the fact that the relevant knowledge for additional rollout of these six vaccines already existed. We know from the more unique but sharper variation coming from the COVID-19 pandemic that entirely new vaccines, together with a new body of scientific knowledge, were created in response to the huge increase in the demand for vaccines against this novel virus (see, e.g., Zuckerman, 2021).

where protection is granted for a fixed term length, creates greater pecuniary incentives for late-stage cancer treatments relative to early-stage treatments and cancer prevention. They show that there is a powerful effect favoring late-stage treatments. This can be interpreted as a difference in markups between two (imperfectly substitutable) treatment modalities targeting the same underlying problem.

Overall, health care and medical technologies provide ample evidence supporting the role of market size in the direction of innovation, and several of the studies show that policy-induced changes in market size have sizable effects on the direction of technology as well. There is additionally some evidence on the role of markups.

4.3 Agriculture

Early work by Hayami and Ruttan (1971) applied ideas from the induced innovation literature to agriculture, focusing on incentives for developing more or less capital-intensive agricultural methods. More recently, Moscona (2022) studied the long-run effects of the soil erosion and reduced soil productivity in the American Midwest following the Dust Bowl and found that agricultural innovation shifted towards more impacted crops, in an apparent effort to make them more productive under the new soil conditions (see also Hornbeck, 2012).

Related work by Moscona and Sastry (2023a) looks at the more recent period and focuses on the changes in environmental conditions caused by global warming. Using granular data on new crops, these authors find that since the middle of the 20th century agricultural innovation has shifted towards crops that have greater exposure to extreme temperatures, and this has been driven by the types of technologies that are most related to environmental adaptation, such as new crop varieties that can be grown in higher temperatures by existing farmers. The innovation response in these two papers is consistent with the predictions of the framework here when price effects are more powerful than market size effects (that is, if $\sigma < 1$). In contrast, if market size effects had been dominant ($\sigma > 1$), innovations should have been redirected towards crops that are less-affected by the Dust Bowl and climate change, and less of the affected crops should have been cultivated. In contrast, it appears that because price effects are stronger, innovation attempted to make up for the reduced productivity of the affected crops.

4.4 Modern Automation Technologies

Following Acemoglu (2003b), Acemoglu and Restrepo (2018, 2021) and Hémous and Olsen (2022), the two factors here can be mapped to capital and labor to capture a reduced-form model of automation. A more microeconomic model of automation and task allocation between capital and labor, as in Zeira (1998), is developed in Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018); see also Autor, Levy and Murnane (2003).

Acemoglu and Restrepo (2022) provide a first empirical study of this issue, exploiting the fact that demographic change is taking place at different rates across countries, and by reducing the industrial workforce, aging is expanding the demand for automation technologies. This paper shows that demographic change has a large impact on the demand for robots and other automation technologies, and then uses patents and exports of intermediate products to establish that countries with aging workforces file more patents for automation technologies and export more intermediates involved in automation. There is no similar impact for non-automation technologies. This study thus establishes a powerful channel from

the market size for automated production methods to both the innovation and adoption of automation technologies. I will use the data from this study in the next section as well.

More recent work by Dechezleprêtre, Hémous, Olsen and Zanella (2022) confirms and more deeply explores this relationship. The authors build a new firm-level data set of automation innovations based on patent text, and combine this with macroeconomic data across 41 countries. They estimate that higher wages for low-skill workers lead to more automation innovation. In addition, they exploit the Hartz labor market reforms in Germany, which led to lower protection for workers, and show that these reforms were associated with a reduction in automation innovations.

Finally, Clemens, Lewis and Postel (2018) study the end of the *Bracero* program, which brought about half a million Mexican immigrants to work in US farming. They find no discernible effect on agricultural wages and provide evidence that this is because the decline in the supply of unskilled labor induced the adoption of more mechanized production methods in US agriculture.

Overall, the evidence suggests that, although many factors have impacted the development and introduction of modern automation techniques, a major boost has come from changes in the market size for these technologies, driven by declines in the supply of labor and corresponding higher wages, due to aging or changes in regulations.

4.5 British Industrial Revolution

There is also a small economic history literature providing evidence that at various turning points during the industrialization process, the direction of innovation was heavily shaped by market sizes and scarcity of labor and other inputs. In addition to Habakkuk and Allen’s work discussed in the Introduction, Hanlon (2015) studies the technological implications of the shortage of cotton in Britain created by the Union Navy’s blockade of Southern shipping during the American Civil War. After the introduction of the cotton gin, the US South had become a major (slave-based) producer of cotton and the largest exporter of this crop to the expanding British industry. The blockade of Southern exports during the Civil War created an acute shortage of inputs to the British cotton industry, which in response turned to alternative cotton varieties, grown in India (and to a lesser extent, in Egypt and Brazil). The spinning technologies used at the time were adapted to the American cotton and could not be used on Indian and other varieties. Hanlon interprets this change as an expansion in the market size of these alternative cotton varieties, which should, according to the framework presented here, trigger a major expansion of complementary technologies. Hanlon documents that this is exactly what happened. There was a flurry of spinning innovations and patenting, but no spike in other textile technologies, such as weaving, and no changes in non-textile patents (for which there was no major change in market size). Moreover, by studying the variation in cotton prices, Hanlon shows that the induced-innovation response was large enough to cause the equivalent of the strong-bias result outlined above.

4.6 Inappropriate Technologies

Another implication of the framework presented here is that when a disproportionate share of innovative activity is concentrated in a few countries, and researchers in these countries target their own economies’ factor endowments and prices, then the global technology will be inappropriate to the needs of remaining countries, especially when their conditions are very different from those of innovative economies.

A recent important paper by Moscona and Sastry (2023b) extends the framework in Acemoglu and Zilibotti (2001) and provides evidence that this inappropriate technology channel is present and quantitatively important. They establish that new crop varieties and seeds are developed to be resistant to pests and pathogens that are important in the US and Western nations, while the major pests and pathogens in the rest of the world, though closely related, are distinct. As a result, the same agricultural technologies do not achieve high productivity in developing-world agriculture. Moscona and Sastry document that inappropriate agricultural technologies are generally not adopted in the developing world and consequently, agricultural productivity remains low in these countries. They estimate that global agricultural output could be increased by about 58%, if the direction of innovation were better targeted towards the agricultural conditions in less developed economies. Relatedly, Acemoglu et al. (2023) show that global medical research responds to disease burden in rich countries, but not in poor countries, and Diao, Ellis, McMillan, and Rodrik (2021) provide evidence from Ethiopia and Tanzania that firms using Western, capital-intensive technologies are not increasing employment.

5 Quantitative Evaluation

In this section, I discuss how the extent of technology distortions can be assessed in the leading applications considered here (automation, health and energy). I first outline the econometric framework I use for estimating the parameters σ and δ , and discuss how markup and externality differences are calibrated. I then provide baseline estimates and a quantitative evaluation of distortions in the direction of innovation in these three sectors.

5.1 Econometric Framework

For automation technologies, I use the data set on automation patents and demographic changes from Acemoglu and Restrepo (2022). For health care, I rely on the medical research and disease burdens data set compiled by Acemoglu et al. (2023). For energy, I use the firm-level patenting and innovation data set constructed by Aghion et al. (2016), who then combine this with information on policy-induced changes in the cost of gasoline.

In each case, I start from the dynamic innovation possibilities frontier for entity (country or firm) f and technology j :

$$\frac{\dot{N}_{fj}(t)}{N_{fj}(t)} = \xi_{fj}(t)\eta_{fj}\Gamma_j(t)N_{fj}(t)^{-\frac{1-\delta}{2}}N_{f\sim j}(t)^{\frac{1-\delta}{2}}S_{fj}(t), \quad (22)$$

which generalizes equation (17) by including a constant, $\eta_{fj} > 0$, parameterizing the productivity of entity f in technology area j , a time effect, $\Gamma_j(t)$, and a random term, $\xi_{fj}(t)$, orthogonal to everything else. In addition, $S_{fj}(t)$ is a measure of research effort devoted by this entity to technology area j (e.g., the number of scientists allocated to this line) and $\delta \in [0, 1)$ again designates the degree of path dependence.

In this formulation, $\dot{N}_{fj}(t)$ is the flow of patents or innovations, while $N_{fj}(t)$ is the stock of patents/innovations, which is estimated following Cockburn and Griliches (1988) and Aghion et al. (2016), by assuming that the stock of knowledge embedded in past patents depreciates at some rate (and as in these papers I set this rate of depreciation to 20%).

When there are only two types of technologies, as in the benchmark model, we can define $n_{ft} \equiv$

$N_{f2}(t)/N_{f1}(t)$ as relative technology, take logs and use the approximations $\Delta n_{ft} \approx \dot{N}_{f2}(t)/\dot{N}_{f1}(t)$ as we transition from continuous to discrete time to obtain:

$$\ln\left(\frac{\Delta n_{ft}}{n_{ft}}\right) = \bar{\eta}_f + \bar{\gamma}_t - \rho \ln n_{ft} + s_{ft} + \bar{\xi}_{ft},$$

where I defined $\bar{\eta}_f \equiv \ln \eta_{fj} - \ln \eta_{f\sim j}$, $s_{ft} \equiv \ln S_{fj}(t) - \ln S_{f\sim j}(t)$, $\bar{\gamma}_t \equiv \ln \Gamma(t) - \ln \Gamma_{\sim j}(t)$, and $\bar{\xi}_{ft} \equiv \ln \xi_{fj}(t) - \ln \xi_{f\sim j}(t)$. I also set $\rho \equiv 1 - \delta$.

Suppose that we have a shifter/forcing variable at the country or firm level z_f (such as relative resource prices, market sizes, or policies) that alters the relative profitability of different technologies. Suppose also that the allocation of research effort between the two technologies at the firm level can be written as $s_{ft} = \chi \ln z_{ft} + \lambda \Delta \ln z_{ft}$.¹⁷ Substituting for this relationship, we arrive at the estimating equation:

$$\ln\left(\frac{\Delta n_{ft}}{n_{ft}}\right) = \bar{\eta}_f + \bar{\gamma}_t - \rho \ln n_{ft} + \chi \ln z_{ft} + \lambda \Delta \ln z_{ft} + \bar{\xi}_{ft}. \quad (23)$$

The left-hand side variable is the flow of relative patents or innovations normalized by stock of relative patents in the two technology areas. The forcing variable is also relative.

From estimates of (23), the key parameters necessary for quantifying the extent of distortions can be recovered. First, I set $\hat{\delta} = \max\{0, 1 - \hat{\rho}\}$, which imposes that the estimate for δ does not become negative in a few specifications in which $\hat{\rho}$ takes a value above one. Moreover, long-run effects can be obtained from estimates of (23). In particular, in an interior BGP we have $\dot{N}_{f2}(t)/N_{f2}(t) = \dot{N}_{f1}(t)/N_{f1}(t)$ in (22) and $s_{ft} = \chi \ln z_{ft}$, and thus the long-run relationship between relative technology and the forcing variable is $\ln n_{ft} = \text{constant} + \frac{\chi}{1-\delta} \ln z_{ft}$. Estimated long-run effects can then be linked to the underlying parameters. Specifically, equation (16) implies that when the forcing variable is changes in market size, we have $\chi/(1-\delta) = (\sigma-1)/(1-\delta\sigma)$, and for the case of energy, where the forcing variable is changes in energy prices, we have $\chi/(1-\delta) = -(\sigma-1)(1-\alpha)/(\alpha\beta(1-\delta\sigma))$.

The same economic relationships can be alternately estimated at the technology-field level, by running the following regression separately by field, which follows directly from (22):

$$\ln\left(\frac{\Delta N_{fjt}}{N_{fjt}}\right) = \eta_{fj} + \Gamma_{jt} - \frac{\rho}{2} \ln N_{fjt} + \chi \ln Z_{fjt} + \lambda \Delta \ln Z_{fjt} + \xi_{fjt}. \quad (24)$$

With only two research areas, this is equivalent to estimating (23). In the medical research regressions, there will be many more than just two areas, and hence focusing on this regression will be more meaningful.

5.2 Measuring Shares, Externalities and Markups

Throughout, I use numbers from the US economy. The factor shares α and β are obtained from the Bureau of Economic Analysis IO Use Tables. Table 1 provides a summary of these numbers for our three applications. For the automation application, I assume $\alpha = 1$ and take β to be the wage bill divided by the sum of the wage bill and expenditures on intermediate inputs for the manufacturing sector in 2012, which gives $\beta = 0.22$. For the health application, I again set $\alpha = 1$ and take β to be the wage bill divided

¹⁷This form follows, for example, when there are within-period diminishing returns or congestion effects in research (e.g., Acemoglu, 1998).

by the sum of the wage bill and expenditures on intermediate inputs for the health care sector in 2012, which gives $\beta = 0.55$. For the energy application, I take $1 - \alpha$ and $\alpha\beta$ to be, respectively, expenditures on material inputs divided by the sum of the wage bill and expenditures on intermediate and material inputs, and the wage bill divided by the sum of the wage bill and expenditures on intermediate and material inputs. This gives $\alpha = 0.86$ and $\beta = 0.32$.

The simplest method to measure the τ parameters is to start with existing estimates of externalities from certain economic activities. I then convert these externalities into the equivalent of $\tilde{\tau}$ in our model, which is in consumption units (recall equation (1)). Throughout I adopt the convention that the sector creating negative externalities is sector 2.

In the automation case, I follow Acemoglu, Manera and Restrepo (2020), who interpret estimates of wage declines following job loss as proxying for quasi-rents that workers enjoy above and beyond the marginal cost of labor hours (and thus above the the social opportunity cost of employment). Hence, if automation technologies reduce employment, they create a negative pecuniary externality proportional to labor earnings. Assuming that for the target group (workers) consumption is approximately equal to labor earnings, this corresponds to $\tilde{\tau}$ in our model. I measure this externality by combining estimates from Acemoglu and Restrepo (2020a) on the effects of robots on employment with the average estimate of the extent of wage declines following job loss (15%, following the review of the literature in Acemoglu, Manera and Restrepo, 2020). The details are provided in Appendix A. The resulting estimate is $\tilde{\tau}_2 = 0.07$, as shown in Table 1. As an alternative, more conservative estimate, I consider the case where only half of the workforce receives quasi-rents, which implies average quasi-rents of 7.5% and $\tilde{\tau}_2 = 0.03$.

In the health care case, I interpret output Y as quality-adjusted life years, which depend on expenditures and innovations in two broad categories: preventative vs. curative technologies (used after the onset of disease). I allow these two types of technologies to have different markups and social benefits. This distinction and my approach are motivated by Kenkel (2000), Kremer and Snyder (2015) and Newhouse (2021). To measure social benefits, I use a sample of 71 new technologies that can be sorted into these two categories and then rely on existing estimates from the medical literature to obtain how much gain in quality-adjusted life years (QALYs) is obtained per one dollar of total cost (upfront R&D spending plus per unit usage costs). These numbers indicate that there are fewer QALY gains from a dollar of spending in curative technologies than preventative technologies, and I interpret this shortfall as a negative externality from N_2 (curative) relative to N_1 (preventative). The baseline estimate of $\tilde{\tau}_2 = 0.37$ indicates that the QALY gains from the preventative category are about 60% larger than those from the curative category. The details of these technologies and the relevant calculations are provided in Appendix C. In the baseline quantitative evaluation for health care, I set these externalities equal to zero and subsequently explore the implications of these additional distortions separately.¹⁸ Broadly speaking, differences in externalities and markups between these two classes of technologies result from the fact that both the level of demand and the elasticity of demand for technologies that can be used after the onset of a disease are different from those for preventative ones because of individual incentives and insurance and public policy reimbursement rules (see Kremer and Snyder, 2015, and Newhouse, 2021).

In the energy application, I focus on negative externalities created by fossil-fuel emissions. I use a

¹⁸Estimating the shortfall of QALYs from curative technologies should be viewed as an alternative to using markup differences, since differential markups will lead to different QALYs from preventative and curative technologies.

worldwide social cost of carbon (CO2) of $SCC = \$185$ dollars per metric ton of carbon (in 2020 dollars), based on Rennert et al.’s (2022) estimate. For the baseline, I focus on US damages only, since the other applications also ignore worldwide externalities. To convert this estimate to US-only damages I use the ratio of US to worldwide damages from the Resources For the Future’s recent report (0.14), which gives $SCC \approx \$26$ per metric ton of carbon.¹⁹ These estimates are then converted into $\tilde{\tau}_2$ following the procedure described in Appendix A. The resulting estimates are depicted in Table 1 as well.

Estimates for $\tilde{\tau}_2$ need to be converted to τ_2 . Recalling that $\tau_j = \tilde{\tau}_j / (\gamma_j^\varepsilon p_j^{1-\varepsilon})$ and also that $\gamma_j^\varepsilon p_j^{1-\varepsilon} = \gamma_j (Y_j/Y)^{\frac{\varepsilon-1}{\varepsilon}} = p_j Y_j/Y$, we have $\tau_2 = \left(\frac{p_2 Y_2}{p_1 Y_1 + p_2 Y_2} \right)^{-1} \tilde{\tau}_2$. Since estimates of $p_1 Y_1$ and $p_2 Y_2$ in the various approaches are likely to be imprecise (because of the difficulty of matching the conceptual categories here to data), I use the fact that this expression implies $\tau_2 \geq \tilde{\tau}_2$ and in the spirit of obtaining lower bounds on innovation distortions, I proxy τ_2 by $\tilde{\tau}_2$ in all three applications. As a result, my baseline estimates of technological externalities are $\tau_2 = 0.07$ in the automation case (or $\tau_2 = 0.03$ using the more conservative estimate of quasi-rents); $\tau_2 = 0.37$ in the health care case; and $\tau_2 = 0.13$ in the energy case when I focus on social cost of carbon for the US and $\tau_2 = 0.94$ when taking full global damages into account.

Finally, I assume that markups are equal between the two technologies in the automation and energy applications. In the health care application, I use data from health-related Compustat firms, sorted into the preventive vs. curative technologies. I then use production function estimation or accounting data to obtain estimates of markups for these two groups of firms.²⁰ The details and list of companies in each category are provided in Appendix C. The baseline markup estimates, which follow De Loecker, Eeckhout and Unger (2020), yield $\mu_1 = 0.46$ and $\mu_2 = 1.70$ for the period 1980-2016, as shown in Table 1.²¹ These markups are high, though broadly consistent with the numbers in De Loecker, et al. (2020). For example, their estimates of revenue-weighted markups for pharmaceutical and medicine manufacturing, and medical equipment and supplies manufacturing (the two four-digit industries most closely related to curative technologies) are, respectively, 3.41 and 2.14 (or cost-weighted markups of 2.97 and 1.91). These high numbers are also in line with the common view that certain medical procedures and pharmaceuticals are priced much above marginal cost in the United States, partly because of lack of regulation and partly because of employer-provided health insurance and Medicare reimbursement policies (see, e.g., Agnell,

¹⁹See <https://www.rff.org/publications/explainers/social-cost-carbon-101/>. Rennert et al.’s (2022) estimate is based on a discount rate of 2%. The EPA’s most recent preferred approach also suggests a similar social cost of carbon (\$190) based on 2% discount rate. See <https://www.epa.gov/system/files/documents/2022-11/>.

²⁰Health-care firms in the preventative category include basic health providers, various companies specialized in diagnostics and vaccine manufacturers, while curative ones include major pharmaceutical companies as well as high-tech medical equipment manufacturers. See Appendix C for a full list.

²¹Markup estimates from Compustat should be interpreted as simply suggestive, since both capital and labor information from this data set are subject to significant measurement error, and it is impossible to separate the output and factor usage of multi-product companies into different business lines. Moreover, as I show in Appendix C, there are nontrivial fluctuations and trends in markup estimates. Nevertheless, Appendix C also shows that using different methods for production function estimation yields quite similar estimates.

One conceptual issue, discussed in Appendix C, is whether markups over marginal cost from variable inputs, as estimated by the production function approach, or accounting markups that subtract payments to quasi-fixed factors are more appropriate in this context. In particular, although accounting profits do not correspond to economic profits, they may be more informative about incentives for innovation and entry. Reassuringly, accounting markups for the two group of firms are comparable to our baseline estimates ($\mu_1 = 0.51$ and $\mu_2 = 1.35$), and using them instead yields broadly similar results, as also shown in Appendix C.

Finally, I experimented with applying the same methods to the energy sector as well, but because there are only a few firms that can be associated with clean technologies, these markups are unstable.

2005, Howard, et al., 2015, Anderson, Hussey and Petrosyan, 2019, and Case and Deaton, 2019).

5.3 Estimates: Automation

In the context of automation, I focus on technologies targeting automation vs. those that can broadly be thought to increase worker productivity. Columns 1 and 2 of Table 2 present estimates of equation (23) using five-year or ten-year patent counts sorted between automation and non-automation technologies. Following Acemoglu and Restrepo (2022), I exploit medium-term, partially anticipated changes in demographics, which reduce the availability of labor to perform manual tasks across countries. I focus on anticipated (15 or 20-year) changes in the ratio of workers aged 56 and above to those between the ages of 25-55 as the measure of aging.²² The left-hand side variable is the relative flow of automation patents compared to the relative stock of automation patents. On the right-hand side, I additionally control for GDP per capita, log population and average years of schooling of the population at the beginning of the sample interacted with time dummies. These controls allow for flexible differential trends as a function of baseline characteristics. As in the original paper, these regressions are weighted by manufacturing employment in 1990, since patent data are significantly noisier for countries with smaller manufacturing employment levels. The sample period in this case is 1986-2015.

Throughout this table, I report heteroscedasticity-robust standard errors clustered to allow for serial-correlation (at the country level in columns 1-4, and at the firm level in columns 5 and 6).

Column 1 in Table 2 depicts estimates from equation (23) for a full sample of 66 countries. The main parameters are estimated reasonably precisely. For example, the estimate of $\hat{\rho} = 0.77$ (standard error = 0.14) implies a value of $\hat{\delta} = 0.23$ for the degree of path dependence. In addition, $\hat{\chi}$ is estimated as 0.87 (standard error= 0.31), which maps to a long-run effect of 1.14—hence, 1% more aging will be associated with 1.14% shifts towards automation technologies. These estimates also imply an elasticity of substitution between factors of $\hat{\sigma} = 1.69$, which ensures that $\hat{\delta}\hat{\sigma} = 0.40 < 1$.

These parameters, together with equation (21), yield a lower bound distortion of $n^{SP}/n^{EQ} = 0.82$, as shown in Panel C at the bottom of the table. This is a sizable difference between the equilibrium and socially optimal direction of technology—a socially-optimal technology ratio that is 18% lower than the equilibrium—despite the fact that the pecuniary externality in the automation case appears small. This magnitude is partly explained by the non-trivial value of $\hat{\delta}\hat{\sigma} = 0.40$, which amplifies the impact of distortions. Nevertheless, the welfare loss from equilibrium distortions is modest, about 1% in consumption-equivalent terms. Panel D shows that using an even smaller estimate of quasi-rents from employment (7.5% instead of 15%) gives correspondingly smaller numbers for the technology distortion ($n^{SP}/n^{EQ} = 0.91$) and welfare losses (0.2%).

Column 2 of Table 2 considers one variation on the automation numbers by using ten-year rather than five-year intervals. The results are broadly similar: $\hat{\rho} = 0.76$ (standard error= 0.12), $\hat{\chi} = 1.16$ (standard error= 0.38) and a long-run effect of 1.52. These imply $\hat{\sigma} = 1.85$ and $\hat{\delta}\hat{\sigma} = 0.44$, which together yield slightly larger technology distortions and welfare costs: $n^{SP}/n^{EQ} = 0.79$ and 1% in Panel C. Panel D numbers are correspondingly smaller.

²²When using five-year (ten-year) changes, the anticipated aging variable is for the next 15 (20) years.

Acemoglu and Restrepo (2021) also show that instrumental-variable estimates exploiting fertility changes from several decades before give very similar results to these ordinary least squares (OLS) estimates. Here I focus on OLS models.

Table B1 in Appendix B presents a number of robustness checks and additional results. In particular, in columns 3-8, I show that similar results hold: when instead of $\ln x$, I use $\ln(1+x)$ and include observations with zeros; when I use the inverse hyperbolic sine, $a \sinh$ (a transformation that allows for zeros and approximately yields logarithmic form for non-zero observations); and for the OECD sample. The implied technology distortion n^{SP}/n^{EQ} remains comparable to those in columns 1 and 2, ranging from 0.56 to 0.76 in Panel C. The exceptions are the inverse hyperbolic sine model and the specification that focuses on just OECD countries, in both cases at the five-year horizon (columns 5 and 7). In these instances, the estimates for δ are higher and consequently technology distortions are more pronounced ($n^{SP}/n^{EQ} = 0.40$ and 0.34) and welfare losses are also larger. Finally, columns 9 and 10 of this table reports estimates of equation (23) from the recent paper by Dechezleprêtre et al. (2022), who study the effects of skill premia on automation technologies at the firm level. Using five-yearly observations across about 1150 firms that have at least four automation patents, these columns show similar estimates of the degree of path dependence and the elasticity of substitution σ to the baseline estimates in columns 1 and 2 of Table 2 (column 9 includes firm fixed effects and industry by time fixed effects, while column 10 additionally includes country by time fixed effects). As a result, we obtain broadly comparable technology distortions using estimates from this firm-level data set: $n^{SP}/n^{EQ} = 0.47$ and 0.61 in the two columns, with welfare losses of 3% and 2%, respectively.

5.4 Estimates: Health

Because detailed data classified into preventative and curative health innovations are not available, for the regression analysis I use data on medical research and disease burdens from Acemoglu et al. (2023). Columns 3 and 4 report estimates from equation (24) using these data. The left-hand side variable is the flow of medical articles for a disease in a country during a particular time period (relative to the stock of medical articles relevant for this observation), and the forcing variable is the disease burden for that disease, country and time. Disease burdens are computed as declines in the number of disability-adjusted life years caused by each disease in a country and time period in our sample.²³ All regressions in this case are unweighted and control for disease, country and time fixed effects.

Column 3 focuses on five-year periods, while column 4 looks at ten-year observations. In both columns, the sample covers the years 1990-2019 and 279 diseases, and comes from 193 countries. In column 3 we have a total of 55,699 observations, while there are 37,389 observations in column 4.

The estimates in the two columns are similar. In column 3, $\hat{\rho}$ is 0.93 (standard error= 0.03), which implies a path dependence parameter of $\hat{\delta} = 0.07$. The estimate of $\hat{\chi} = 0.10$ combined with these numbers yields a long-run effect of 0.11. Hence, a 1% increase in the burden of a specific disease in a country leads to a 0.11% increase in the medical research directed to that disease. The implied elasticity of substitution is $\hat{\sigma} = 1.10$, which again puts us comfortably in the region where $\hat{\delta}\hat{\sigma} = 0.08 < 1$. In column 4, ρ is estimated to be a little more than 1 (1.11), which implies no path dependence and thus I set $\delta = 0$. Other estimates remain similar: in particular, a long-run effect of 0.14 and $\sigma = 1.14$.

Panel C focuses on markup differences between preventative and curative categories, given in Table 1 (and ignores differences in externalities). Technology distortions are similar in the two columns:

²³These calculations are based on data from the Global Burden of Disease (GBD) project, which is a collaboration between the World Bank and the Institute for Health Metrics and Evaluation (IHME). See Acemoglu et al. (2023) for details.

$n^{SP}/n^{EQ} = 0.43$ in column 3 and 0.45 in column 4, meaning that the technology ratio is about 45% biased in favor of curative technologies in the decentralized equilibrium. The resulting welfare effects are sizable—around 5% or 6% (which should be interpreted as a fraction of health care consumption).

Panel D looks at the implications of the τ_2 estimate from the shortfall of QALY gains from curative technologies relative to preventative technologies (now ignoring markup differences). This alternative way of conceptualizing misaligned innovation incentives in health care leads to even larger technology distortions: n^{SP}/n^{EQ} is around 0.6, and welfare losses from the equilibrium direction of technology are correspondingly bigger (17-18%).

Table B2 in Appendix B, I show that the estimates reported in columns 3 and 4 of Table 2, and thus the implied technology distortions and welfare effects, are quite robust. Similar results are obtained when instead of $\ln x$, I use $\ln(1+x)$ and keep observations with zeros; when I use the inverse hyperbolic sine ($a \sinh$) transformation; when the country fixed effects are omitted; when we include country times year and disease times year fixed effects; and when we focus only in variation in the United States. The implied values for n^{SP}/n^{EQ} in Panel C are mostly around 0.4 and the welfare effects are also comparable to those in Table 2, except in specifications using $\ln(1+x)$ and $a \sinh$ with five-year observations and in the two specifications that do not include country fixed effects, where technology distortions are larger, ranging between 0.11 and 0.17, and the welfare effects are correspondingly more substantial.

5.5 Estimates: Energy

In the context of energy, I follow the conceptual structure in Acemoglu et al. (2012) that distinguishes dirty (coal, gas and oil) technologies and clean (renewables and nuclear) technologies. Columns 5 and 6 of Table 2 use data from Aghion et al. (2016) and report firm-level regressions of the flow of patents of clean technologies relative to dirty technologies in the automobile sector, once again normalized with their respective stocks. In these data, there are many observations with zero stocks, and I follow Aghion et al. and include these observations by using $\ln(1+x)$. This gives a data set of 3,412 firms across 58 countries for the years 1986-2005, and 13,684 and 6,824 observations in the two columns.

The estimates are fairly similar between column 5 and 6. For example, ρ is estimated as 0.81 (standard error= 0.03) in column 5 and 0.86 (standard error= 0.04) in column 6. Long-run effects are comparable as well: -1.89 in column 5 and -1.23 in column 6 (these are the effects of higher gasoline prices, leading to lower clean technology patents, hence the negative sign). The estimated values of σ are also similar across the two columns: 2.73 and 2.53. As a result, in both columns, we have $\hat{\delta}\hat{\sigma} < 1$.

Using our baseline estimate of $\tau_2 = 0.13$ in Panel C of Table 1 based on social cost of carbon for the United States, the technology distortion is found to be $n^{SP}/n^{EQ} = 0.44$ in column 5, and a little smaller, $n^{SP}/n^{EQ} = 0.57$ in column 6. These are again sizable distortions with welfare losses of about 2-3%.

Instead, with global damages we have $\tau_2 = 0.94$, and because this externality is close to one, equation (20) implies that the social planner would like to essentially shut down fossil-fuel technologies (i.e. $n^{SP}/n^{EQ} \approx 0$), as indicated in Panel D.

In Appendix B, I report various robustness checks. The general pattern is broadly comparable to that shown in columns 5 and 6. Long-run effects and elasticity estimates are quite similar, including in specifications that add spillovers from the stock of innovation of other firms in the same country, as in Aghion et al. (2016). The extent of technology distortions, n^{SP}/n^{EQ} , remains fairly stable and ranges

between 0.37 and 0.74 across all specifications in Panel C.

Overall, in all three of the applications considered here, I find suggestive evidence that distortions in the direction of technology can be sizable and generate non-trivial welfare consequences. These results should be interpreted with ample caution, since both the estimates of the underlying parameters and even more so the estimates of externalities and markups are subject to considerable uncertainty. They are presented in the spirit of suggestive evidence to stimulate more work in this area.

6 Concluding Remarks

Technological change is vital for continued economic prosperity and can help tackle many of the epochal challenges facing humanity, such as climate change, pandemics and global poverty. Because of its society-wide benefits, corporations and individuals tend to underinvest in innovation, and this underinvestment provides a central justification for government support for science, academia and corporate R&D. But will the “market process”—working through profit incentives, competition and reputational concerns of researchers—get the direction of innovation right?

Typically, there are many alternative technologies and paradigms even within a narrow field. In health care, innovation can be directed towards curative technologies and pharmaceuticals, or it can prioritize preventative technologies. In energy and transport, innovation can be directed towards clean or dirty alternatives. In most industries, researchers and corporations decide how much to invest to automate production processes vs. how much to prioritize increasing worker marginal productivity, by providing better tools, new labor-intensive tasks and new learning opportunities to employees. In agriculture, novel crop varieties can target pests and pathogens that are pervasive in some countries ahead of others.

In this paper, I have suggested that there may be systemic reasons for the direction of innovation to be distorted. Using a simple framework, I highlighted the factors impacting the direction of technology, and illustrated how economic or social externalities (such as carbon emissions) and markup differences between technologies can lead to a misaligned direction of innovation. Innovation distortions tend to reduce or even reverse welfare gains from technological progress (for example, when research effort focuses on socially costly technologies) and can even slow down economic growth (for example, because of markup differences).

There are three distinct objections one could raise to the approach in this paper. *First*, even if the market does not get the direction of innovation completely right, governments and bureaucrats could be worse at it. This objection is valid and is the reason why much of my discussion focused on systemic sources of distortions that can be determined without superior technical knowledge on the part of bureaucrats or some impressive ability to “pick winners”. If there are markup differences across the products generated by different technologies or quantifiable externalities—as I have proposed—the extent of distortions can be determined and agreed upon.

Second, one may argue that distortions resulting from the direction of technology are secondary relative to underinvestment in overall innovation and/or they are small relative to other costs that government intervention in the innovation process would generate. This is also a valid concern, but ultimately the extent of these distortions is a quantitative question. For this reason, I provided evidence from three distinct domains on distortions in the direction of technology.

Third, attempts to deal with distortions in the direction of innovation could lead to new and challenging political economy questions. I return to this important question at the end of these remarks.

In light of these caveats, the current paper should be seen as a first step in a more detailed investigation of possible distortions in the direction of technological change and potential remedies. This is the reason why the theoretical framework is chosen to be as simple as possible and the quantitative evaluation is purely suggestive. Several interesting questions are open for future study within this framework, and I list some of them here.

- It would be instructive to model and empirically investigate the extent to which other social factors can also create distortions in the direction of scientific and corporate research. One possibility is researchers following each other's leads and becoming influenced by each other's visions to such an extent that makes them overinvest in some paradigms. I have suggested in past work (Acemoglu and Restrepo, 2020b, and Acemoglu and Johnson, 2023) that this may be a concern within the field of artificial intelligence, pushing researchers to prioritize automation and mass-scale data collection. What the theoretical microfoundations of such effects are, whether this type of bias is indeed present in practice, and whether governmental or societal intervention may be possible in this case are interesting questions for future research.
- The theoretical analysis in the paper ignored the interplay between Pigovian taxes and policy aimed at redirecting technological change. A critical question both from a theoretical and an applied point of view is to what extent these different classes of policies are complements or substitutes.
- Much industrial policy became mired in corruption and political problems in the past, and one may be worried that any government intervention aimed at influencing the direction of technological change would be similarly hampered by political economy challenges. This is particularly true since history is full of examples of special interest groups attempting to block technological change to protect their rents or privileges (e.g., Acemoglu and Robinson, 2012). On the other hand, the endogeneity of the direction of innovation opens up new political economy avenues, and studying them is a fruitful area for future inquiry (see Acemoglu and Johnson, 2023).
- In this context, another research area is to model the market structure of the relevant industries in greater detail, so that the pro- or anti-competitive effects of policies aimed at redirecting technological change can be evaluated. For example: can firms and researchers be encouraged to invest in socially more beneficial technologies without reducing the extent of competition in the economy?
- Lastly, the empirical part of the current paper was a first attempt, and more systematic work on measuring distortions in the direction of innovation is a critical area for future research.

References

- Acemoglu, Daron (1998) “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *Quarterly Journal of Economics*, 113(4): 1055-1090.
- Acemoglu, Daron (2002) “Directed Technical Change.” *The Review of Economic Studies*, 69(4): 781–809.
- Acemoglu, Daron (2003a) “Patterns of Skill Premia.” *The Review of Economic Studies*, 70(2): 199–230.
- Acemoglu, Daron (2003b) “Labor- and Capital-Augmenting Technical Change.” *Journal of the European Economic Association*, 1(1): 1–37.
- Acemoglu, Daron (2007) “Equilibrium Bias of Technology.” *Econometrica*, 75(5): 1371-1409.
- Acemoglu, Daron (2010) “Diversity and Technological Progress.” in *The Rate and Direction of Inventive Activity Revisited*, ed. Josh Lerner and Scott Stern (Chicago, University of Chicago Press, 2010), 319-356.
- Acemoglu, Daron (2010) “When Does Labor Scarcity Encourage Innovation?” *Journal of Political Economy*, 118(6): 1037–1078.
- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hémous (2012) “The Environment and Directed Technical Change.” *American Economic Review*, 102(1): 131-66.
- Acemoglu, Daron, Nicholas Ajzeman, Cevat Giray Aksoy, Martin Fiszbein, and Carlos Molina (2021) “(Successful) Democracies Read Their Own Support.” NBER Working Paper No. 29167.
- Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr (2016) “Transition to Clean Technology.” *Journal of Political Economy*, 124(1): 52-140.
- Acemoglu, Daron, Ufuk Akcigit, and William Kerr (2016) “Innovation Network.” *Proceedings of the National Academy of Sciences*, 113(41): 11483-11488.
- Acemoglu, Daron, and David Autor (2012) “Skills, Tasks and Technologies: Implications for Employment and Earnings.” in *Handbook of Labor Economics: Volume 4B*, ed. David Card, Orley Ashenfelter (Amsterdam, Elsevier, 2011), 1043-1171.
- Acemoglu, Daron, David Cutler, Amy Finkelstein, and Joshua Linn (2006) “Did Medicare Induce Pharmaceutical Innovation?” *American Economic Review*, 96(2): 103-107.
- Acemoglu, Daron, David Hémous, Lint Barrage and Philippe Aghion (2019) “Climate Change, Directed Innovation, and Energy Transition: The Long-Run Consequences of the Shale Gas Revolution.” 2019 Meeting Paper #1302, Society for Economic Dynamics.
- Acemoglu, Daron, and Simon Johnson (2023) *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*. Public Affairs, New York.
- Acemoglu, Daron, Michael I. Jordan and Glen Weyl (2021) “The Turing Test is Bad for Business.” wired.com. November 8, 2021. <https://www.wired.com/story/artificial-intelligence-turing-test-economics-business/>.
- Acemoglu, Daron, and Todd Lensman (2023) “Technology Paradigms, Lock-in, and Economic Growth.” in progress.
- Acemoglu, Daron, and Joshua Linn (2004) “Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry.” *Quarterly Journal of Economics*, 119(3): 1049–1090.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020) “Does the US Tax Code Favor Automation?” *Brookings Papers on Economic Activity*, (Spring): 231–285.

Acemoglu, Daron, Moscona, Jacob, Karthik A. Sastry and Heidi Williams (2023) “The Rich-World Bias of Medical Research” work in progress.

Acemoglu, Daron, and Pascual Restrepo (2018) “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review*, 108(6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo (2020a) “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy*, 128(6): 2188–2244.

Acemoglu, Daron, and Pascual Restrepo (2020b) “The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand.” *Cambridge Journal of Regions, Economy and Society*, 13(1): 25–35.

Acemoglu, Daron, and Pascual Restrepo (2022) “Demographics and Automation.” *The Review of Economic Studies*, 89(1): 1–44.

Acemoglu, Daron, and James A. Robinson (2012) *Why Nations Fail: The Origins of Power, Prosperity and Poverty*. Crown Business, New York.

Acemoglu, Daron, David Y. Yang, and Jie Zhou (2023) “Power and the Direction of Research: Evidence from China’s Academia.” MIT Working Paper.

Acemoglu, Daron, and Fabrizio Zilibotti (2001) “Productivity Differences.” *Quarterly Journal of Economics*, 116(2): 563–606.

Aghion, Philippe, Roland Bénabou, Ralf Martin, and Alexandra Roulet (2020) “Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?” National Bureau of Economic Research Working Paper #26921.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous and John Van Reenen (2016) “Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry.” *Journal of Political Economy*, 124(1): 1–51.

Agnell, Marcia (2004) *The Truth About Drug Companies: How They Deceive Us and What to Do About It*. Random House, New York.

Ahmad, Syed (1966) “On the Theory of Induced Invention.” *The Economic Journal*, 76(302): 344–357.

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde (2021) “Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth.” *The Review of Economic Studies*, 88(1): 1–43.

Allen, Robert C. (2009) *The British Industrial Revolution in Global Perspective*. Cambridge University Press, Cambridge.

Anderson, Gerard F., Peter Hussey and Varduhi Petrosyan (2019) “It’s Still The Prices, Stupid: Why The US Spends So Much On Health Care, And A Attribute To Uwe Reinhardt.” *Health Affairs*, 38(1): 87–95.

Arthur, W. Brian (1989) “Competing Technologies, Increasing Returns, and Lock-In by Historical Events.” *The Economic Journal*, 99(394): 116–31.

Autor, David H., Frank Levy and Richard J. Murnane (2003) “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 118(4): 1279–1333.

Autor, David H., David Dorn, and Gordon Hanson (2019) “When Work Disappears: How Adverse La-

bor Market Shocks Affect Fertility, Marriage, and Children’s Living Circumstances.” *American Economic Review: Insights*, 1(2): 161-178

Azoulay, Pierre, Joshua S. Graff Zivin, Danielle Li, and Bhaven N. Sampat (2019) “Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules.” *The Review of Economic Studies*, 86(1): 117–152.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen (2013) “Identifying Technology Spillovers and Product Market Rivalry.” *Econometrica*, 81(4): 1347-1393.

Bovenberg, A. Lans and Sjak Smulders (1995) “Environmental Quality and Pollution-Augmenting Technological Change in a Two-Sector Endogenous Growth Model.” *Journal of Public Economics*, 57(3): 369-391.

Branstetter, Lee G., Guangwei Li and Mengjia Ren (2022) “Picking Winners? Government Subsidies and Firm Productivity in China.” National Bureau of Economic Research Working Paper #30699.

Brynjolfsson, Erik (2022) “The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence.” *Dædalus*, Spring 2002, 272-287.

Budish, Eric, Benjamin N. Roin, and Heidi Williams (2015) “Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials.” *American Economic Review*, 105(7): 2044-85.

Case, Anne and Angus Deaton (2020). *Deaths of Despair and the Future of Capitalism*. Princeton University Press, Princeton New Jersey.

Clemens, Michael A., Ethan G. Lewis, and Hannah M. Postel (2018) “Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion.” *American Economic Review*, 108(6): 1468-87.

Cockburn, Iain, and Zvi Griliches (1988) “Industry Effects and Appropriability Measures in the Stock Market’s Valuation of R&D and Patents.” *American Economic Review*, 78(2): 419-423.

Coe, David T., and Elhanan Helpman (1995) “International R&D Spillovers.” *European Economic Review*, 39(5): 859-887.

Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams (2019) “The More We Die, The More We Sell? A Simple Test of the Home-Market Effect.” *Quarterly Journal of Economics*, 134(2): 843–894.

Dechezleprêtre, Antoine, David Hémous, Morten Olsen and Carlo Zanella (2021) “Induced Automation: Evidence from Firm-Level Patent Data” (unpublished manuscript, April 2021).

Dechezleprêtre, Antoine, Ralf Martin and Myra Mohnen (2014) “Knowledge Spillovers from Clean and Dirty Technologies.” London School of Economics, Centre for Economic Performance Discussion Paper #1300.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020) “The Rise of Market Power and the Macroeconomic Implications.” *Quarterly Journal of Economics*, 135(2): 561–644.

Di Maria, Corrado, and Simone Valente (2008) “Hicks Meets Hotelling: The Direction of Technical Change in Capital–Resource Economies.” *Environment and Development Economics*, 13(6): 691–717.

Diao, Xinshen, Mia Ellis, Margaret S. McMillan, and Dani Rodrik (2021) “Africa’s Manufacturing Puzzle: Evidence from Ethiopian and Tanzanian Firms.” NBER Working Paper No. 28344.

Dosi, Giovanni (1982) “Technological Paradigms and Technological Trajectories: A Suggested Interpretation of the Determinants and Directions of Technical Change.” *Research Policy*, 11(3): 147-162.

- Drandakis, E. M., and E. S. Phelps. “A Model of Induced Invention, Growth and Distribution.” *The Economic Journal*, 76(304): 823–840.
- Finkelstein, Amy (2004) “Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry.” *Quarterly Journal of Economics*, 119(2): 527–564.
- Gancia, Gino, and Fabrizio Zilibotti (2005) “Horizontal Innovation in the Theory of Growth and Development,” in *Handbook of Economic Growth: Volume 1A*, ed. Philippe Aghion, Steven N. Durlauf (Amsterdam, Elsevier, 2005), 111-170.
- Goulder, Lawrence H., and Stephen H. Schneider (1999) “Induced Technological Change and the Attractiveness of CO2 Abatement Policies.” *Resource and Energy Economics*, 21(3–4): 211-253.
- Grimaud, André, and Luc Rouge (2008) “Environment, Directed Technical Change and Economic Policy.” *Environmental and Resource Economics*, 41: 439–463.
- Grossman, Gene M., and Elhanan Helpman (1993). *Innovation and Growth in the Global Economy*. MIT Press, Cambridge.
- Gruber, Jonathan, and Simon Johnson (2019) *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*. Hachette, UK.
- Habakkuk, H. J. (1962) *American and British Technology in the Nineteenth Century: Search for Labor Saving Inventions*. Cambridge University Press, Cambridge.
- Hager, Thomas (2009) *The Alchemy of Air: a Jewish Genius, a Doomed Tycoon, and the Scientific Discovery that Fed the World but Fueled the Rise of Hitler*. Broadway Books, New York.
- Hanlon, W. Walker (2015) “Necessity Is the Mother of Invention: Input Supplies and Directed Technical Change.” *Econometrica*, 83(1): 67-100.
- Hayami, Yujiro, and V. W. Ruttan (1970) “Factor Prices and Technical Change in Agricultural Development: The United States and Japan, 1880-1960.” *Journal of Political Economy*, 78(5): 1115–1141.
- Hémous, David (2016) “The Dynamic Impact of Unilateral Environmental Policies.” *Journal of International Economics*, 103(November 2016): 80-95.
- Hémous, David, and Morten Olsen (2022) “The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality.” *American Economic Journal: Macroeconomics*, 14(1): 179-223.
- Hicks, John (1932) *The Theory of Wages*. Macmillan, London.
- Hornbeck, Richard (2012) “The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe.” *American Economic Review*, 102(4): 1477-1507.
- Howard, David H., Peter B. Bach, Ernst R. Berndt and Rena M. Conti (2015) “Pricing in the Market for Anticancer Drugs.” *Journal of Economic Perspectives*, 29(1): 139-62.
- Howell, Sabrina T., Jason Rathje, John Van Reenen and Jun Wong (2021) “Opening up Military Innovation: Causal Effects of ‘Bottom-up’ Reforms to U.S. Defense Research.” London School of Economics, Programme on Innovation and Diffusion Working Paper #4.
- Hufbauer, Gary Clyde, and Euijin Jung (2021) *Scoring 50 Years of US Industrial Policy, 1970–2020*. Peterson Institute for International Economics, Washington, DC.
- Jones, Charles I. (2005) “The Shape of Production Functions and the Direction of Technical Change.” *Quarterly Journal of Economics*, 120(2): 517–549.
- Jones, Charles I. and John C. Williams (1998) “Measuring the Social Return to R&D.” *Quarterly*

Journal of Economics, 113(4): 1119-1135.

Kennedy, Charles (1964) "Induced Bias in Innovation and the Theory of Distribution." *The Economic Journal*, 74(295): 541-547.

Kenkel, Donald S. (2000) "Prevention." in *Handbook of Health Economics: Volume 1*, ed. Anthony J. Culyer, Joseph P. Newhouse (Amsterdam, Elsevier, 2000), 1675-1720.

Kiley, M. T. (1999) "The Supply of Skilled Labour and Skill-biased Technological Progress." *The Economic Journal*, 109(458): 708-724.

Koyama, Mark, and Jared Rubin (2022) *How the World Became Rich: The Historical Origins of Economic Growth*. Polity Press, Cambridge.

Kremer, Michael, and Christopher M. Snyder (2015) "Preventives Versus Treatments." *Quarterly Journal of Economics*, 130(3): 1167-1239.

Lane, Nathan (2021) "Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea." University of Oxford, Center for the Study of African Economies Working Paper #2021-09-02.

Martin, Ralf and Dennis Verhoeven (2022) "Knowledge Spillovers from Clean and Emerging Technologies in the UK." Paper presented at International Association for Research in Income and Wealth 37th General Conference, Luxembourg, August 22-26, 2022: IARIW, <https://iariw.org/wp-content/uploads/2022/08/Martin-Verhoeven-IARIW-2022.pdf>.

Mazzucato, Mariana (2015) *Entrepreneurial State: Debunking Public vs. Private Sector Myths*. Anthem Press, New York.

Mitrunen, Matti (2019) "War Reparations, Structural Change, and Intergenerational Mobility." Institute for International Economic Studies, Stockholm University, Working Paper.

Mokyr, Joel (1992) *The Levers of Riches: Technological Creativity and Economic Progress*. Oxford University Press, New York.

Mokyr, Joel (2011) *The Gifts of Athena*. Princeton University Press, Princeton.

Moretti, Enrico, Claudia Steinwender and John Van Reenen (2020) "The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers." London School of Economics, Centre for Economic Performance Discussion Paper #1662.

Moscona, Jacob (2021) "Environmental Catastrophe and the Direction of Invention: Evidence from the American Dust Bowl" (unpublished manuscript, August 24, 2021).

Moscona, Jacob, and Karthik A. Sastry (2023a) "Does Directed Innovation Mitigate Climate Damage? Evidence from U.S. Agriculture." *Quarterly Journal of Economics*, forthcoming.

Moscona, Jacob, and Karthik A. Sastry (2023b) "Inappropriate Technology: Evidence from Global Agriculture." (unpublished manuscript, November 15, 2022).

Neumann, Peter J., and Joshua T. Cohen. (2009) "Cost Savings and Cost-Effectiveness of Clinical Preventive Care." *The Synthesis Project*, Research Synthesis Report Sep(18), <https://europepmc.org/article/med/22052182>.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1999) "The Induced Innovation Hypothesis and Energy-Saving Technological Change." *Quarterly Journal of Economics*, 114(3): 941-75.

Newhouse, Joseph P. (2021) "An Ounce of Prevention." *Journal of Economic Perspectives*, 35(2): 101-118.

Pack, Howard, and Kamal Saggi (2006) "Is There a Case for Industrial Policy? A Critical Survey."

The World Bank Research Observer, 21(2): 267–297.

Phillips, Kathryn A., and David R. Holtgrave (1997) “Using Cost-Effectiveness/Cost-Benefit Analysis to Allocate Health Resources: A Level Playing Field for Prevention?” *American Journal of Preventive Medicine*, 13(1): 18-25.

Popp, David (2002) “Induced Innovation and Energy Prices.” *American Economic Review*, 92(1): 160-180.

Rennert, Kevin et al. (2022) “Comprehensive Evidence Implies a Higher Social Cost of CO₂.” *Nature*, 610(7933): 687-692

Rodrik, Dani (2014) “Green Industrial Policy.” *Oxford Review of Economic Policy*, 30(3): 469-491.

Ridley, Matt (2020) *How Innovation Works: And Why It Flourishes in Freedom*. Harper, New York.

Samuelson, Paul A. (1965) “A Theory of Induced Innovation along Kennedy-Weisäcker Lines.” *The Review of Economics and Statistics*, 47(4): 343–356.

Thoenig, Mathias, and Thierry Verdier (2003) “A Theory of Defensive Skill-Biased Innovation and Globalization.” *American Economic Review*, 93(3): 709-728.

Wilson, William Julius (1996) *When Work Disappears: The World of the New Urban Poor*. Alfred Knopf, New York.

Zeira, Joseph (1998) “Workers, Machines, and Economic Growth.” *Quarterly Journal of Economics*, 113(4): 1091–1117.

Zuckerman, Gregory (2021) *A Shot to Save the World: The Inside Story of the Life-or-Death Race for a COVID-19*. Portfolio, UK.

Tables

Table 1: Externally Calibrated Parameters

Parameters	Description	Values
Panel A: Automation		
α	1 – Material Share	1
β	Labor Share divided by α	0.22
(μ_1, μ_2)	Markups (assumption)	(μ, μ)
(τ_1, τ_2)	Externality (quasi-rent = 15%)	(0,0.07)
(τ_1, τ_2)	Externality (quasi-rent = 7.5%)	(0,0.03)
Panel B: Health		
α	1 – Material Share	1
β	Labor Share divided by α	0.55
(μ_1, μ_2)	Markups (estimated)	(0.46,1.70)
(τ_1, τ_2)	Externality (from QALYs)	(0,0.37)
Panel C: Energy		
α	1 – Material Share	0.86
β	Labor Share divided by α	0.32
(μ_1, μ_2)	Markups (assumption)	(μ, μ)
(τ_1, τ_2)	Externality (US damages)	(0,0.13)
(τ_1, τ_2)	Externality (World damages)	(0,0.94)

Notes: This table presents the values of the parameters used in the equilibrium and welfare analysis. Panel A is for the automation application, Panel B for the health care application, and Panel C for the energy application. Material and labor shares are taken from the BEA Use Table for 2012 (see text for details). Markups in Panel B are computed from Compustat via the production function estimation method based on De Loecker et al. (2020). Firm-level markups are aggregated to the technology level using firm cost shares. Appendix C provides more details and alternative estimates. Externalities are computed from wage declines following job loss, based on Acemoglu, Manera and Restrepo (2020) in Panel A; from the shortfall of quality-adjusted life year gains from curative technologies relative to preventative technologies (based on own calculations in Appendix C) in Panel B; and from Rennert et al.’s (2022) estimate of the social cost of CO₂, converted to US-equivalent damages and for world-wide damages in Panel C. Further details are provided in the text, Appendix A and Appendix C.

Table 2: Estimates and Implied Parameters

Application	Automation		Health		Energy	
Frequency	5-year	10-year	5-year	10-year	5-year	10-year
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A Parameters Estimated from Regressions</i>						
Initial Relative Stock: $\hat{\rho}$	0.77 (0.14)	0.76 (0.12)	0.93 (0.03)	1.11 (0.03)	0.81 (0.03)	0.86 (0.04)
Initial Shifter: $\hat{\chi}$	0.87 (0.31)	1.16 (0.38)	0.10 (0.01)	0.14 (0.01)	-1.52 (0.29)	-1.06 (0.66)
Changes in Shifter: $\hat{\lambda}$	1.18 (0.43)	1.81 (0.63)	-0.004 (0.02)	0.001 (0.02)	-0.45 (0.20)	1.12 (0.82)
Observations	232	125	55,699	37,389	13,648	6,824
<i>Panel B Implied Parameters</i>						
Long-run Effects	1.14	1.52	0.11	0.14	-1.89	-1.23
$\hat{\delta}$	0.23	0.24	0.07	0.00	0.19	0.14
$\hat{\sigma}$	1.69	1.85	1.10	1.14	2.73	2.53
$\hat{\varepsilon}$	4.09	4.82	1.18	1.26	7.27	6.56
$\hat{\delta}\hat{\sigma}$	0.40	0.44	0.08	0.00	0.53	0.36
<i>Panel C Equilibrium and Welfare Comparison</i>						
n^{SP}/n^{EQ}	0.82	0.79	0.43	0.45	0.44	0.57
$U^{SP} - U^{EQ}$	0.01	0.01	0.06	0.05	0.03	0.02
<i>Panel D Equilibrium and Welfare Comparison(Alternatives)</i>						
n^{SP}/n^{EQ}	0.91	0.89	0.58	0.59	0.00	0.00
$U^{SP} - U^{EQ}$	0.002	0.002	0.18	0.17	13.74	8.94

Notes: This table presents regression estimates (Panel A), implied parameter values (Panel B) and implied distortions and welfare results (Panels C and D) for the three applications. In all cases, regressions are estimated with ordinary least squares and heteroscedasticity-robust standard errors are presented in parentheses. Standard errors are clustered at country level in columns 1-4 and at the firm-level in columns 5-6. Odd-numbered columns are for five-year changes and even-numbered columns are for ten-year changes. Columns 1 and 2 are for the automation application and are at the country-time period level and present regressions weighted by manufacturing employment in 1990. The dependent variable is the number of newly granted patents for automation technologies relative to other utility patents divided by the stock of patents related to automation relative to the stock of other utility patents (in logs). Shifters are the level and change in the ratio of workers above the age of 56 to workers between 21 and 55. Column 1 uses expected 20-year change, and column 2 uses expected a 15-year change (in logs). Both columns include region dummies, and the 1990 values of log GDP per capita, log of population, average years of schooling and the ratio of workers above 56 to workers aged 21 in 1990 interacted with period dummies. Columns 3 and 4 are for the health care application, and observations are at the country-disease-period level. The dependent variable is relative number of new medical articles for each disease divided by relative stock of medical articles for that disease (in logs). Shifters are the log of the burden of disease for the relevant country-year-disease cell. Both columns include country, disease and period fixed effects. Columns 5 and 6 are for the energy application and observations are at the firm-period level. The dependent variable is relative number of newly granted patents for dirty technologies relative to newly granted patents for clean technologies (with the log (1+x) transformation). Shifters are firm-level fuel prices adjusted (based on firm-level fuel consumption) inclusive of taxes. Both columns include firm and period fixed effects as well as the values of government R&D subsidies for clean innovation, regulations over emissions, the relevant country's GDP per capita for that period (as in Aghion et al., 2016). In columns 1 and 2, Panel C uses 15% quasi-rent for workers, and Panel D uses 7.5% quasi-rents. In columns 3 and 4, Panel C focuses on markup differences and Panel D uses the externality estimate computed from the shortfall of quality-adjusted life years from curative relative to preventative technologies. In columns 5 and 6, Panels C and D use externality numbers based on Rennert et al.'s (2022) estimate of the social cost of CO2 for the US and the world, respectively.

ONLINE APPENDICES FOR DISTORTED INNOVATION: DOES THE MARKET GET THE DIRECTION OF TECHNOLOGY RIGHT?

Appendix A: Additional Theoretical Results and Omitted Proofs

Derivation of Static Equilibrium

In this part of the Appendix, I provide a few expressions omitted from the text. First, the maximization of (9) gives the demand for machine varieties and resource inputs as

$$x_j(\nu) = \left[p_j \left(\frac{(1-\beta)\alpha}{(1+\mu_j)\psi} \right)^\alpha \left(\frac{1-\alpha}{q_j^R} \right)^{1-\alpha} \right]^{\frac{1}{\alpha\beta}} L_j, \quad (\text{A1})$$

and

$$R_j = \left[p_j \left(\frac{(1-\beta)\alpha}{(1+\mu_j)\psi} \right)^{\alpha(1-\beta)} \left(\frac{1-\alpha}{q_j^R} \right)^{1-\alpha+\alpha\beta} \right]^{\frac{1}{\alpha\beta}} N_j L_j. \quad (\text{A2})$$

Substituting these into (4), we obtain the levels of sectoral production as

$$Y_j = \left(\frac{(1-\beta)\alpha}{(1+\mu_j)\psi} \right)^{\frac{1-\beta}{\beta}} \left(\frac{1-\alpha}{q_j^R} \right)^{\frac{1-\alpha}{\alpha\beta}} p_j^{\frac{1-\alpha\beta}{\alpha\beta}} N_j L_j. \quad (\text{A3})$$

Combining this expression for $j = 1, 2$ with (11) and rearranging yields (12) in the text.

I next show that the equilibrium characterized in the text is unique when $\delta\sigma < 1$ and there are always multiple (corner) equilibria when $\delta\sigma > 1$. Recall from footnote 9 that each scientist has a mass $\mathfrak{s} > 0$, and then we are taking the limit case where $\mathfrak{s} \rightarrow 0$. Then consider an allocation in which all researchers work in sector 2 (of course, the argument is analogous when they all work in sector 1). For this allocation to be an equilibrium, we need that switching to sector 1 is not profitable for an individual scientist. This requires

$$\tilde{\eta}_2 (\bar{S} - \mathfrak{s})^{\frac{\delta}{1-\delta}} \pi_2 \geq \tilde{\eta}_1 \mathfrak{s}^{\frac{\delta}{1-\delta}} \pi_1.$$

Rearranging this equation and using the equivalent conditions from the main text, it becomes

$$\frac{\tilde{\eta}_2}{\tilde{\eta}_1} \left(\frac{\bar{S} - \mathfrak{s}}{\mathfrak{s}} \right)^{\frac{\delta}{1-\delta}} \frac{\mu_2}{\mu_1} \left(\frac{1+\mu_2}{1+\mu_1} \right)^{-\frac{1}{\beta}} \left(\frac{q_2^R}{q_1^R} \right)^{-\frac{1-\alpha}{\alpha\beta}} p^{\frac{1}{\alpha\beta}} \left(\frac{L_2}{L_1} \right) \geq 1,$$

or substituting from (12) and recalling that $N_j = \tilde{\eta}_j S_j^{\frac{1}{1-\delta}}$, it is equivalent to

$$\frac{\eta_2}{\eta_1} \left(\frac{\gamma_2}{\gamma_1} \right)^{\frac{\varepsilon}{\sigma}} \frac{\mu_2}{\mu_1} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{-\frac{\sigma - (1-\beta)}{\beta\sigma}} \left(\frac{q_2^R}{q_1^R} \right)^{-\frac{(\sigma-1)(1-\alpha)}{\alpha\beta\sigma}} \left(\frac{L_2}{L_1} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{\bar{S} - \mathfrak{s}}{\mathfrak{s}} \right)^{-\frac{1-\delta\sigma}{(1-\delta)\sigma}} \geq 1.$$

Now taking the limit $\mathfrak{s} \rightarrow 0$, we can see that this condition can never be satisfied when $\delta\sigma < 1$, since $\left(\frac{\bar{S} - \mathfrak{s}}{\mathfrak{s}} \right)^{-\frac{1-\delta\sigma}{(1-\delta)\sigma}} \rightarrow 0$, and hence the entire left-hand side limits to 0. Conversely, when $\delta\sigma > 1$, $\left(\frac{\bar{S} - \mathfrak{s}}{\mathfrak{s}} \right)^{-\frac{1-\delta\sigma}{(1-\delta)\sigma}} \rightarrow +\infty$ and thus the left-hand side limits to $+\infty$, ensuring that this condition is always satisfied as a strict inequality. This establishes that when $\delta\sigma < 1$, there are no corner equilibria and the interior equilibrium characterized in the text is unique. Conversely, when $\delta\sigma > 1$, corner allocations are always equilibria.

Derivation of Socially-Optimal Technology Ratio in the Static Model

The first-order conditions for the social planner in the static model can be written as:

$$\tilde{\eta}_1 S_1^{\frac{\delta}{1-\delta}} \left[\frac{d \ln Y}{d N_1} + \frac{d \ln E}{d N_1} \right] = \tilde{\eta}_2 S_2^{\frac{\delta}{1-\delta}} \left[\frac{d \ln Y}{d N_2} + \frac{d \ln E}{d N_2} \right].$$

From this expression, using the fact that $d \ln N_j = d N_j / N_j$ and substituting N_j for S_j from $N_j = \tilde{\eta}_j S_j^{\frac{1}{1-\delta}}$ (with $\eta_j \equiv \tilde{\eta}_j^{1-\delta}$), we get (19) in the main text.

Next, note that

$$\frac{d \ln Y}{d \ln Y_j} = \gamma_j \left(\frac{Y_j}{Y} \right)^{\frac{\varepsilon-1}{\varepsilon}} = \gamma_j^\varepsilon p_j^{1-\varepsilon},$$

where the second relationship exploits the representative household's utility maximization condition (recalling that the social planner does not directly manipulate prices). Moreover:

$$\begin{aligned} \frac{d \ln Y_2}{d \ln N_2} &= 1 + \frac{\partial \ln Y_2}{\partial \ln p_2} \frac{d \ln p_2}{d \ln N_2} \\ &= 1 + \frac{1 - \alpha\beta}{\alpha\beta} \frac{d \ln p_2}{d \ln p} \frac{d \ln p}{d \ln N_2} \end{aligned}$$

$$\begin{aligned} \frac{d \ln Y_2}{d \ln N_1} &= \frac{\partial \ln Y_2}{\partial \ln p_2} \frac{d \ln p_2}{d \ln N_1} \\ &= \frac{1 - \alpha\beta}{\alpha\beta} \frac{d \ln p_2}{d \ln p} \frac{d \ln p}{d \ln N_1} \end{aligned}$$

$$\begin{aligned} \frac{d \ln Y_1}{d \ln N_1} &= 1 + \frac{\partial \ln Y_1}{\partial \ln p_1} \frac{d \ln p_1}{d \ln N_1} \\ &= 1 + \frac{1 - \alpha\beta}{\alpha\beta} \frac{d \ln p_1}{d \ln p} \frac{d \ln p}{d \ln N_1} \end{aligned}$$

$$\begin{aligned}\frac{d \ln Y_1}{d \ln N_2} &= \frac{\partial \ln Y_1}{\partial \ln p_1} \frac{d \ln p_1}{d \ln N_2} \\ &= \frac{1 - \alpha\beta}{\alpha\beta} \frac{d \ln p_1}{d \ln p} \frac{d \ln p}{d \ln N_2}\end{aligned}$$

Finally, using the ideal price condition, (13),

$$\begin{aligned}\frac{dp_1/dp}{p_1} &= -\gamma_2^\varepsilon p^{-\varepsilon} p_1^{1-\varepsilon} \\ \frac{dp_2/dp}{p_2} &= \gamma_1^\varepsilon p^{\varepsilon-2} p_2^{1-\varepsilon}.\end{aligned}$$

Or

$$d \ln p_1/d \ln p = -\gamma_2^\varepsilon p_2^{1-\varepsilon} \text{ and } d \ln p_2/d \ln p = \gamma_1^\varepsilon p_1^{1-\varepsilon}.$$

Now combining these expressions, we have

$$\begin{aligned}\frac{d \ln Y}{d \ln N_1} &= \gamma_1^\varepsilon p_1^{1-\varepsilon} \left[1 - \frac{(1 - \alpha\beta)\gamma_2^\varepsilon p_2^{1-\varepsilon}}{\sigma} \right] + \gamma_2^\varepsilon p_2^{1-\varepsilon} \left[\frac{(1 - \alpha\beta)\gamma_1^\varepsilon p_1^{1-\varepsilon}}{\sigma} \right] \\ &= \gamma_1^\varepsilon p_1^{1-\varepsilon},\end{aligned}$$

and

$$\begin{aligned}\frac{d \ln Y}{d \ln N_2} &= \gamma_1^\varepsilon p_1^{1-\varepsilon} \left[\frac{(1 - \alpha\beta)\gamma_2^\varepsilon p_2^{1-\varepsilon}}{\sigma} \right] + \gamma_2^\varepsilon p_2^{1-\varepsilon} \left[1 - \frac{(1 - \alpha\beta)\gamma_1^\varepsilon p_1^{1-\varepsilon}}{\sigma} \right] \\ &= \gamma_2^\varepsilon p_2^{1-\varepsilon}.\end{aligned}$$

Combining these expressions, we obtain the desired result:

$$\frac{d \ln Y}{d \ln N_j} = \gamma_j^\varepsilon p_j^{1-\varepsilon} \text{ for } j = 1, 2.$$

Finally, using the same steps as in the previous subsection of the Appendix, we can show that when $\delta\sigma < 1$, the second-order conditions of the social planner's optimization problem are always satisfied in the interior allocation given by (16). Conversely, when $\delta\sigma > 1$, the interior allocation is not a local maximum, and instead there are two local maxima at the corners, with all scientists working in one of the two sectors. One of these two local maxima is the global maximum. Which one is preferred can be easily determined by using the expression for welfare derived in the next subsection of the Appendix and substituting for S_j in terms of N_j (once again from $N_j = \tilde{\eta}_j S_j^{\frac{1}{1-\delta}}$), and comparing the resulting expressions as $\varepsilon \rightarrow 0$.

Measuring Externalities

In the theoretical analysis, I simplified the discussion by assuming that externalities are created directly by technology choices. This means that I need to convert existing externality estimates into those that appear in the form of the $\tilde{\tau}_j$ or τ_j variables. I now discuss how this can be done.

Automation: In the automation case, I follow Acemoglu, Manera and Restrepo's (2020) review of

the literature. The median estimate of quasi-rents (and thus pecuniary externalities) in labor income is about 15%. I combine this with Acemoglu and Restrepo’s (2020) estimate of the effect of robot adoption on the employment in local labor markets (approximated by commuting zones in the US). Namely, let us equate automation technologies with N_2 , and employment with L_1 (and L_2 can be capital or high-skilled labor working with automated technologies), and denote the working age population by Pop. Then we have

$$\begin{aligned}
\tilde{\tau}_2 &= -\frac{d \ln E}{d \ln N_2} = -\frac{d \ln E}{d \ln L_1} \frac{d \ln L_1}{d \ln N_2} \\
&= -\frac{d \ln E}{d \ln L_1} \frac{d \ln L_1}{d L_1} \frac{d L_1}{d L_1 / \text{Pop}} \frac{d L_1 / \text{Pop}}{d N_2} \frac{d N_2}{d \ln N_2} \\
&= -\frac{d \ln E}{d \ln L_1} \frac{\text{Pop}}{L_1} \frac{d L_1 / \text{Pop}}{d N_2} N_2 \\
&= -(0.15) \times \frac{1}{0.63} \times (-0.39) \times 0.73 \\
&= 0.07,
\end{aligned}$$

where -0.15 is from Acemoglu, Manera and Restrepo’s review of the literature, 0.63 is the employment to population ratio in the United States, averaged over the years 1990-2007 in Current Population Survey,²⁴ -0.39 is Acemoglu and Restrepo’s (2020) estimate of the impact of one more robot per 1000 industrial workers on employment to population ratio, and 0.73 is their estimate of the stock of robots between 1993 and 2007. This number implies that a proportional increase in automation technology creates a 7% negative pecuniary externality on workers. I then convert this into τ_2 as described above.

Health care: In the health care case, the main distortion I focus on is differential markups, which does not need any conversion. Secondly, I compute the externalities in terms of differences in quality-adjusted life year returns per one dollar of spending on technology between the preventative and high-tech/late-stage curative technologies (inclusive of R&D costs and usage costs). These numbers are therefore directly comparable to $\tilde{\tau}_j$ ’s in our model. Further details of medical procedures, drugs and technologies used in these computations and the studies from which the estimates are taken are provided in Appendix C.

Energy: In the energy case, I use estimates of the social cost of carbon. In this framework, carbon corresponds to the (suitably rescaled) resource input R_2 (identifying dirty technologies with sector 2). The social cost of carbon is in terms of the impact of one more metric ton of carbon emissions on consumption-equivalent welfare. The externality in the utility equation (1) is in terms of proportional effect on consumption. Therefore, I compute $\tilde{\tau}_2$ as follows:

$$\begin{aligned}
\tilde{\tau}_2 &= -\frac{d \ln E_2}{d \ln N_2} = -\frac{d E_2}{d R_2} \frac{R_2}{E_2} \frac{d \ln R_2}{d \ln N_2} \\
&= SCC \times \frac{\text{CO2 emission}}{\text{Energy Consumption}} \frac{d \ln R_2}{d \ln N_2},
\end{aligned}$$

where I am using the fact that the relevant consumption is total energy consumption and proxying $d \ln R_2 / d \ln N_2 \simeq 1$.

²⁴From FRED, <https://fred.stlouisfed.org/series/EMRATIO>

Welfare Comparisons

The welfare difference between the social optimum in the equilibrium can be written as

$$U^{SP} - U^{EQ} = \ln Y(n^{SP}) - \ln Y(n^{EQ}) + \ln E(n^{SP}) - \ln E(n^{EQ}),$$

where I am using the fact that all other endogenous variables are functions of n . The basic idea is to develop the approximation:

$$\Delta \ln Y^{EQ,SP} \equiv \ln Y(n^{SP}) - \ln Y(n^{EQ}) \simeq \frac{d \ln Y(n^{EQ})}{dS} [S^{SP} - S^{EQ}], \quad (\text{A4})$$

where S^{SP} is the allocation of scientists consistent with a technology ratio of n^{SP} , and S^{EQ} is the allocation of scientists implied by the technology ratio of n^{EQ} .

To do this, consider the impact of a change in the allocation of scientists from the equilibrium n^{EQ} , and let $S_1 = S$ and $S_2 = \bar{S} - S$. Then we can write:

$$\begin{aligned} \frac{d \ln Y}{dS} &= \frac{d \ln Y}{d \ln N_1} \frac{d \ln N_1}{dS} - \frac{d \ln Y}{d \ln N_2} \frac{d \ln N_2}{dS} \\ &= \frac{d \ln Y}{d \ln N_1} \frac{dN_1}{dS} \frac{1}{N_1} - \frac{d \ln Y}{d \ln N_2} \frac{dN_2}{dS} \frac{1}{N_2} \\ &= \frac{1}{1-\delta} \left[\gamma_1^\varepsilon p_1^{1-\varepsilon} \tilde{\eta}_1 N_1^{-1} S_1^{\frac{\delta}{1-\delta}} - \gamma_2^\varepsilon p_2^{1-\varepsilon} \tilde{\eta}_2 N_2^{-1} S_2^{\frac{\delta}{1-\delta}} \right] \\ &= \frac{1}{1-\delta} \left[\gamma_1^\varepsilon p_1^{1-\varepsilon} \eta_1 N_1^{-(1-\delta)} - \gamma_2^\varepsilon p_2^{1-\varepsilon} \eta_2 N_2^{-(1-\delta)} \right] \\ &= \frac{1}{1-\delta} \left[\eta_1 \gamma_1^\varepsilon [\gamma_1^\varepsilon + \gamma_2^\varepsilon p^{1-\varepsilon}]^{-1} N_1^{-(1-\delta)} - \eta_2 \gamma_2^\varepsilon [\gamma_1^\varepsilon + \gamma_2^\varepsilon p^{1-\varepsilon}]^{-1} p^{1-\varepsilon} N_2^{-(1-\delta)} \right] \\ &= \frac{N_1^{-(1-\delta)}}{1-\delta} \left[\eta_1 \gamma_1^\varepsilon [\gamma_1^\varepsilon + \gamma_2^\varepsilon (p)^{1-\varepsilon}]^{-1} - \eta_2 \gamma_2^\varepsilon [\gamma_1^\varepsilon + \gamma_2^\varepsilon (p)^{1-\varepsilon}]^{-1} p^{1-\varepsilon} n^{-(1-\delta)} \right] \quad (\text{A5}) \end{aligned}$$

Here, the third line simply uses the expressions for dN_j/dS_j from the static innovation possibilities frontier (5), while the fourth line uses the same transformation as in the text: $N_j = \tilde{\eta}_j S_j^{\frac{1}{1-\delta}}$ and $\eta_j \equiv \tilde{\eta}_j^{1-\delta}$. The penultimate line uses the ideal price condition (13) to substitute p_1 and p_2 in terms of the relative price p . The final line simply factors out N_1 .

Hence,

$$\begin{aligned}
\Delta \ln Y^{EQ,SP} &\simeq \frac{d \ln Y}{dS} (S^{SP} - S^{EQ}) \\
&= \frac{d \ln Y}{dS} \left(\frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{SP})^{1-\delta}} - \frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta}} \right) \\
&= \frac{(N_1^{EQ})^{-(1-\delta)}}{1-\delta} \left(\begin{array}{c} \eta_1 \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} \\ -\eta_2 \gamma_2^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} (p^{EQ})^{1-\varepsilon} (n^{EQ})^{-(1-\delta)} \end{array} \right) (S^{SP} - S^{EQ}) \\
&= \frac{(N_1^{EQ})^{-(1-\delta)}}{1-\delta} \eta_1 \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} \left(1 - \frac{\eta_2}{\eta_1} \left(\frac{\gamma_2}{\gamma_1} \right)^\varepsilon (p^{EQ})^{1-\varepsilon} (n^{EQ})^{-(1-\delta)} \right) (S^{SP} - S^{EQ}) \\
&= \frac{(N_1^{EQ})^{-(1-\delta)}}{1-\delta} \eta_1 \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} \left[1 - \frac{\mu_1}{\mu_2} \frac{1 + \mu_2}{1 + \mu_1} \right] (S^{SP} - S^{EQ}) \\
&= \frac{(N_1^{EQ})^{-(1-\delta)}}{1-\delta} \eta_1 \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} \left[1 - \frac{\mu_1}{\mu_2} \frac{1 + \mu_2}{1 + \mu_1} \right] \left(\frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{SP})^{1-\delta}} - \frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta}} \right) \\
&= \frac{1}{1-\delta} \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1} \left[1 - \frac{\mu_1}{\mu_2} \frac{1 + \mu_2}{1 + \mu_1} \right] \left(\frac{1 + \frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta}}{1 + \frac{\eta_1}{\eta_2} (n^{SP})^{1-\delta}} - 1 \right).
\end{aligned}$$

In these derivations, I have used the following steps. The second line is from (A4), while the third line substitutes from (A5). The fourth line factors out $\eta_1 \gamma_1^\varepsilon (\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon})^{-1}$. The fifth line uses the expressions for p^{EQ} and n^{EQ} from (12) and (16). The sixth line uses the fact that from (5), the equilibrium and socially-optimal allocations of scientists have to satisfy

$$S^{EQ} = \frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta}} \text{ and } S^{SP} = \frac{\bar{S}}{1 + \frac{\eta_1}{\eta_2} (n^{SP})^{1-\delta}}.$$

The seventh line then substitutes for

$$N_1^{EQ} = \left(\frac{\bar{S}}{\frac{1}{\eta_1} + \frac{(n^{EQ})^{1-\delta}}{\eta_2}} \right)^{\frac{1}{1-\delta}}, \quad (\text{A6})$$

and cancels out terms.

Here everything is a function of n^{EQ} and parameters.

With no markup differences, it can be verified that $\Delta \ln Y^{EQ,SP} = 0$ (which also follows from an application of the envelope theorem). Note, in particular, that the terms in square brackets are equal to

zero. Therefore, with no markup differences (as in the automation and energy cases), we have

$$\begin{aligned} U^{SP} - U^{EQ} &\simeq \ln E(n^{SP}) - \ln E(n^{EQ}) \\ &= (\tilde{\tau}_1 + \tilde{\tau}_2) \ln \left(\frac{N_1^{EQ}}{N_1^{SP}} \right) + \tilde{\tau}_2 \ln \left(\frac{n^{EQ}}{n^{SP}} \right). \end{aligned}$$

This can be computed given the expression for N_1^{EQ} in (A6), and its analogue for the socially-optimal level of technology:

$$N_1^{SP} = \left(\frac{\bar{S}}{\frac{1}{\eta_1} + \frac{(n^{SP})^{1-\delta}}{\eta_2}} \right)^{\frac{1}{1-\delta}}.$$

Hence,

$$\frac{N_1^{EQ}}{N_1^{SP}} = \left(\frac{1 + \frac{\eta_1}{\eta_2} (n^{SP})^{1-\delta}}{1 + \frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta}} \right)^{\frac{1}{1-\delta}}.$$

To proxy for the ratio N_1^{EQ}/N_1^{SP} we only need an estimate of η_2/η_1 . This ratio can be obtained by using the expressions for (12), (16), and (A3). Combining these equations, we obtain:

$$\frac{\eta_2}{\eta_1} = (n^{EQ})^{1-\delta} \left(\frac{\mu_2}{\mu_1} \right)^{-1} \left(\frac{1 + \mu_2}{1 + \mu_1} \right) \left(\frac{p_2 Y_2}{p_1 Y_1} \right)^{-1},$$

where $p_2 Y_2/p_1 Y_1$ is the relative output of the two sectors.

When there are markup differences and no externalities (as in our baseline health care application), then

$$U^{SP} - U^{EQ} \simeq \Delta \ln Y^{EQ,SP}$$

as derived above. This can be computed if we can also compute p^{EQ} and have an estimate for γ_2/γ_1 . In the health care application, we have $\alpha = 1$ and there is no specialized labor, so the same health care labor forces allocated between the two technologies, which implies

$$\frac{w_2^{EQ}}{w_1^{EQ}} = (p^{EQ})^{\frac{1}{\beta}} \left(\frac{1 + \mu_2}{1 + \mu_1} \right)^{-\frac{1-\beta}{\beta}} (n^{EQ})^{\frac{\sigma-1}{\sigma}} = 1.$$

Given markups and n^{EQ} , this equation gives p^{EQ} . To obtain an estimate for γ_2/γ_1 , note first that $\gamma_1 + \gamma_2 = 1$, and thus $\gamma_2/\gamma_1 = (1 - \gamma_1)/\gamma_1$. Moreover,

$$p = \frac{1 - \gamma_1}{\gamma_1} \left(\frac{Y_2}{Y_1} \right)^{-\frac{1}{\varepsilon}},$$

which can be rearranged to yield

$$\gamma_1 = \left[1 + (p^{EQ})^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{p_2 Y_2}{p_1 Y_1} \right)^{\frac{1}{\varepsilon}} \right]^{-1},$$

which again uses an estimate of the relative output of the two sectors.²⁵

This discussion clarifies that to compute the welfare losses from innovation distortions we need two more numbers in each applications: n^{EQ} and $(p_2Y_2)/(p_1Y_1)$. We use the following estimates for these quantities:

- *Automation*: n^{EQ} is taken as the ratio of the total number of automation patents to total non-automation patents across all countries in 2005 from Acemoglu and Restrepo (2022). This gives $n^{EQ} = 0.15$. I also set $(p_2Y_2)/(p_1Y_1) = 0.38$ on the basis of the model-based inference in Acemoglu, Manera and Restrepo (2020), which yields that about 28% of tasks are automated in the US economy.
- *Health care*: n^{EQ} is taken as the ratio of the sum of the discounted stock of curative patents to that of the sum of preventative patents across countries, which gives $n^{EQ} = 16.2$. I proxy p_1Y_1 as total spending on ambulatory health services and social assistance in 2020 from the U.S. Census Bureau Service Annual Survey (SAS) and total spending on diagnostic substances and biological product manufacturing in 2020 from the U.S. Census Bureau and Annual Survey of Manufactures (ASM). I set p_2Y_2 equal to the 2020 revenues of the same industries classified as curative in Appendix C. These revenues are also taken from the ASM. The resulting ratio is $(p_2Y_2)/(p_1Y_1) = 0.13$.²⁶
- *Energy*: $n^{EQ} = 2.20$ is taken as the sum of the stock of dirty patents to that of the stock of clean patents across all countries in 2005, from Aghion et al. (2016). In addition, $(p_2Y_2)/(p_1Y_1)$ is proxied by the ratio of the revenue of renewable energy to that of non-renewable energy, where revenue is calculated as the product of average wholesale electricity price of an energy source and its primary energy consumption from the EIA Monthly Energy Review. The resulting ratio is $(p_2Y_2)/(p_1Y_1) = 3.08$.²⁷

²⁵An alternative approximation for welfare in this case can be derived by taking a first-order Taylor approximation in terms of deviations between technology ratios and then substituting out some of the technology terms by using the same equilibrium relationship we used in the welfare computations, in particular, $\frac{\eta_1}{\eta_2} (n^{EQ})^{1-\delta} = \left(\frac{\gamma_2}{\gamma_1}\right)^\varepsilon (p^{EQ})^{1-\varepsilon} \left[1 - \frac{\mu_1}{\mu_2} \frac{1+\mu_2}{1+\mu_1}\right]^{-1}$. This gives the following approximation for output differences between the optimal and equilibrium allocations in the presence of markup differences:

$$\gamma_1^\varepsilon \left(\gamma_1^\varepsilon + \gamma_2^\varepsilon (p^{EQ})^{1-\varepsilon}\right)^{-1} \left[1 - \frac{\mu_1}{\mu_2} \frac{1+\mu_2}{1+\mu_1}\right] \left(\frac{\left(\frac{\gamma_2}{\gamma_1}\right)^\varepsilon (p^{EQ})^{1-\varepsilon} \left[1 - \frac{\mu_1}{\mu_2} \frac{1+\mu_2}{1+\mu_1}\right]^{-1}}{1 + \left(\frac{\gamma_2}{\gamma_1}\right)^\varepsilon (p^{EQ})^{1-\varepsilon} \left[1 - \frac{\mu_1}{\mu_2} \frac{1+\mu_2}{1+\mu_1}\right]^{-1}}\right).$$

This approximation removes the need to separately estimate η_2/η_1 . In practice, the two expressions give very similar estimates of welfare costs of distorted technology in the health care case.

²⁶More specifically, the preventative categories are: NAICS 621 (Ambulatory health), NAICS 624 (Social assistance), NAICS 325413 (In-vitro diagnostic substances manufacturing), and NAICS 325414 (Biological product manufacturing). The curative categories are: NAICS 325412 (Pharmaceutical preparation manufacturing), NAICS 334510 (Electromedical and electrotherapeutic apparatus manufacturing), NAICS 339112 (Surgical and medical instrument manufacturing), and NAICS 339113 (Surgical appliance and supplies manufacturing). See Appendix C for details. The SAS and ASM data can be accessed at www.census.gov/programs-surveys/sas/data/tables.html and www.census.gov/programs-surveys/asm/data/tables.html, respectively.

²⁷Wholesale prices are from the United States Energy Information Agency (EIA) Power Operations Report (see www.eia.gov/energyexplained/us-energy-facts/) and energy consumption data are from the EIA Monthly Energy Review (www.eia.gov/todayinenergy/detail.php?id=45436).

Dynamic Model

In this part of the Appendix, I provide a few more details about the dynamic framework provided in the text. First recall that when a scientist invents a new machine for sector $j \in \{1, 2\}$, she receives the net present discounted value of future profits from the sale of this machine, given by

$$V_j(t) = \int_t^\infty e^{-\int_t^{t'} r(t'') dt''} \pi_j(t') dt', \quad (\text{A7})$$

where $r(t)$ is the market interest rate at time t , and $\pi_j(t)$ is the common profit that all machines for sector $j \in \{1, 2\}$ will make at time t .

The representative household's optimization problem implies that the growth rate of consumption has to satisfy

$$\frac{\dot{C}(t)}{C(t)} = r(t) - \rho, \quad (\text{A8})$$

as well as a standard transversality condition, which requires the net present discounted value of current and future machine varieties to be finite (see Acemoglu, 2002).

In BGP, consumption has to grow at a constant rate, and thus the interest rate will be constant. Therefore, we have

$$V_j = \frac{\pi_j}{r} \text{ for } j = 1, 2.$$

Using these expressions for the two sectors and combining them with the equilibrium allocation of scientists, we obtain (16), as claimed in the text, which also establishes Proposition 2. The proof for Proposition 3 follows the analysis in Acemoglu (2002) closely and I do not present it here to avoid repetition.

I next consider the socially optimal choice of technology in this dynamic framework. Once again, assuming that the social planner only controls the allocation of scientists, this problem can be written as

$$\max_{\{S(t), N_1(t), N_2(t)\}_0^\infty} \int_0^\infty e^{-\rho t} U[N_1(t), N_2(t)] dt$$

subject to

$$\dot{N}_1(t) = \eta_1 N_1(t)^{(1+\delta)/2} N_2(t)^{(1-\delta)/2} S(t) \quad (\text{A9})$$

and

$$\dot{N}_2(t) = \eta_2 N_1(t)^{(1-\delta)/2} N_2(t)^{(1+\delta)/2} [\bar{S} - S(t)]. \quad (\text{A10})$$

Here $U[N_1(t), N_2(t)] = \ln C(t) + \ln E(t)$ is the level of utility at time t , inclusive of externalities, given the vector of technologies (state variables), $N_1(t)$ and $N_2(t)$. This expression exploits the fact that the level of final good production and hence consumption only depend on the current state of technologies. (All other endogenous variables, and in particular prices of the intermediates, $p_1(t)$ and $p_2(t)$, are solved out as in the equilibrium allocation in the text).

Suppressing time dependence when this will cause no confusion and assigning co-state variables λ_1 and λ_2 to (A9) and (A10), the necessary condition from the maximum principle applied to this optimal

control problem yields:

$$\lambda_1 \eta_1 N_1^{\frac{1+\delta}{2}} N_2^{\frac{1-\delta}{2}} - \lambda_2 \eta_2 N_1^{\frac{1-\delta}{2}} N_2^{\frac{1+\delta}{2}} \begin{cases} > 0 & \implies S = \bar{S} \\ = 0 & \implies S \in [0, \bar{S}] \\ < 0 & \implies S = 0 \end{cases} . \quad (\text{A11})$$

Therefore, just like in the equilibrium, the social planner's solution leads to a bang-bang solution. Moreover, for an interior BGP, we need scientists to be assigned to both sectors, and thus this expression should be equal to zero. Hence, in an interior BGP, we must have:

$$\lambda_1 \eta_1 = \lambda_2 \eta_2 n^\delta . \quad (\text{A12})$$

In order to characterize the socially-optimal technology choices, we need to know the values and evolution of the co-state variables, which are given by the following two differential equations:

$$\begin{aligned} \varrho \lambda_1 - \dot{\lambda}_1 &= \frac{dU}{dN_1} + \frac{1+\delta}{2} \lambda_1 \eta_1 \left(\frac{N_2}{N_1} \right)^{\frac{1-\delta}{2}} S + \frac{1-\delta}{2} \lambda_2 \eta_2 \left(\frac{N_2}{N_1} \right)^{\frac{1+\delta}{2}} [\bar{S} - S] \\ &= \frac{dU}{dN_1} + \frac{1+\delta}{2} \lambda_1 \eta_1 n^{\frac{1-\delta}{2}} S + \frac{1-\delta}{2} \lambda_2 \eta_2 n^{\frac{1+\delta}{2}} (\bar{S} - S) \\ &= \frac{dU}{dN_1} + \lambda_1 \eta_1 n^{\frac{1-\delta}{2}} \left[\frac{1+\delta}{2} S + \frac{1-\delta}{2} (\bar{S} - S) \right], \\ \varrho \lambda_2 - \dot{\lambda}_2 &= \frac{dU}{dN_2} + \frac{1-\delta}{2} \lambda_1 \eta_1 \left(\frac{N_2}{N_1} \right)^{-\frac{1+\delta}{2}} S + \frac{1+\delta}{2} \lambda_2 \eta_2 \left(\frac{N_2}{N_1} \right)^{-\frac{1-\delta}{2}} [\bar{S} - S] \\ &= \frac{dU}{dN_2} + \frac{1-\delta}{2} \lambda_1 \eta_1 n^{-\frac{1+\delta}{2}} S + \frac{1+\delta}{2} \lambda_2 \eta_2 n^{-\frac{1-\delta}{2}} (\bar{S} - S) \\ &= \frac{dU}{dN_2} + \lambda_2 \eta_2 n^{-\frac{1-\delta}{2}} \left[\frac{1-\delta}{2} S + \frac{1+\delta}{2} (\bar{S} - S) \right]. \end{aligned}$$

In BGP, we need $\dot{\lambda}_1 = \dot{\lambda}_2 = 0$, and hence

$$\begin{aligned} \lambda_1 &= \frac{1}{\varrho} \left(\frac{dU}{dN_1} + \lambda_1 \eta_1 n^{\frac{1-\delta}{2}} \left[\frac{1+\delta}{2} S + \frac{1-\delta}{2} (\bar{S} - S) \right] \right) \\ &= \frac{\frac{dU}{dN_1}}{\varrho - \eta_1 n^{\frac{1-\delta}{2}} \left[\frac{1+\delta}{2} S + \frac{1-\delta}{2} (\bar{S} - S) \right]}, \text{ and} \\ \lambda_2 &= \frac{1}{\varrho} \left(\frac{dU}{dN_2} + \lambda_2 \eta_2 n^{-\frac{1-\delta}{2}} \left[\frac{1-\delta}{2} S + \frac{1+\delta}{2} (\bar{S} - S) \right] \right) \\ &= \frac{\frac{dU}{dN_2}}{\varrho - \eta_2 n^{-\frac{1-\delta}{2}} \left[\frac{1-\delta}{2} S + \frac{1+\delta}{2} (\bar{S} - S) \right]}. \end{aligned}$$

Moreover, the scientist allocation has to satisfy the BGP condition:

$$\frac{S}{\bar{S} - S} = \frac{\eta_2}{\eta_1} n^{-(1-\delta)}. \quad (\text{A13})$$

Substituting the values of the co-state variables into (A12), we obtain

$$\eta_1 \frac{\frac{d \ln Y}{d \ln N_1} + \frac{d \ln E}{d \ln N_1}}{\varrho - \eta_1 n^{\frac{1-\delta}{2}} \left[\frac{1+\delta}{2} S + \frac{1-\delta}{2} (\bar{S} - S) \right]} = \eta_2 n^{-(1-\delta)} \frac{\frac{d \ln Y}{d \ln N_2} + \frac{d \ln E}{d \ln N_2}}{\varrho - \eta_2 n^{-\frac{1-\delta}{2}} \left[\frac{1-\delta}{2} S + \frac{1+\delta}{2} (\bar{S} - S) \right]}. \quad (\text{A14})$$

This condition is different from (19) because the social planner takes into account the knowledge spillovers the two sectors create, which have differential effects depending on the relative technology ratio. In the special case where $\delta = 1$, we can combine (A13) and (A14) to show that these differential knowledge spillovers cancel out and we end up with the following condition for the socially-optimal BGP technology ratio:

$$\eta_1 \left[\frac{d \ln Y}{d \ln N_1} + \frac{d \ln E}{d \ln N_1} \right] = \eta_2 \left[\frac{d \ln Y}{d \ln N_2} + \frac{d \ln E}{d \ln N_2} \right],$$

which is identical to (20) when $\delta = 1$, and thus the same n^{SP} in (20) in the text characterizes the socially-optimal BGP technology ratio. The general case where $\delta < 1$ captures the same economic forces I emphasized in the text, but does not admit a closed-form solution for the technology ratio.

Appendix B: Robustness Checks

This part of the Appendix provides robustness checks on the regression results reported in Table 2 in the text. Table B1 considers variations for the automation regressions, Table B2 presents robustness checks for the regressions on the relationship between medical research and disease burden, and finally Table B3 focuses on the relationship between fuel prices and direction of innovation in automobiles. The results of all three tables are discussed in the text.

The formulae for the path dependence parameter δ and the elasticity of substitution σ in the various tables and columns are:

Table B1, columns 1-8 (Long-run effects from relative market sizes)

$$\hat{\delta} = \max \{0, 1 - \hat{\rho}\}, \text{ and}$$

$$\hat{\sigma} = \frac{1 + \hat{\chi} - \hat{\delta}}{1 + \hat{\delta}\hat{\chi} - \hat{\delta}}.$$

Table B1, columns 9-10 (Long-run effects from relative prices and with spillovers)

$$\hat{\delta} = \max \{0, 1 - \hat{\rho} - \hat{\rho}_{\text{spillover}}\}, \text{ and}$$

$$\hat{\sigma} = \frac{2\hat{\chi} + \hat{\delta} - 1 - \hat{\chi}\hat{\delta}}{\hat{\chi} + \hat{\delta} - 1},$$

where $\hat{\rho}_{\text{spillover}}$ is the coefficient on the (relative technology) spillover term.

Table B2, columns 1-12 (Long-run effects from relative market sizes)

$$\hat{\delta} = \max \{0, 1 - \hat{\rho}\}, \text{ and}$$

$$\hat{\sigma} = \frac{1 + \hat{\chi} - \hat{\delta}}{1 + \hat{\delta}\hat{\chi} - \hat{\delta}}.$$

Table B3, columns 1-10 (Long-run effects from relative input prices)

$$\hat{\delta} = \max \{0, 1 - \hat{\rho}\}, \text{ and}$$

$$\hat{\sigma} = \frac{\alpha\beta\hat{\chi} - (1 - \hat{\delta})(1 - \alpha)}{\alpha\beta\hat{\delta}\hat{\chi} - (1 - \hat{\delta})(1 - \alpha)}.$$

Table B3, columns 11-12 (Long-run effects from relative input prices and with spillovers)

$$\hat{\delta} = \max \{0, 1 - \hat{\rho} - \hat{\rho}_{\text{spillover}}\}, \text{ and}$$

$$\hat{\sigma} = \frac{\alpha\beta\hat{\chi} - (1 - \hat{\delta})(1 - \alpha)}{\alpha\beta\hat{\delta}\hat{\chi} - (1 - \hat{\delta})(1 - \alpha)}.$$

Table B1: Robustness for Automation Application

LHS	$\ln(x)$	$\ln(x)$	$\ln(1+x)$	$\ln(1+x)$	$\operatorname{asinh}(x)$	$\operatorname{asinh}(x)$	$\ln(x)$	$\ln(x)$	$\ln(1+x)$	$\ln(1+x)$
Samples	Full	Full	Full	Full	Full	Full	OECD	OECD	Firm-Level	Firm-Level
Frequency	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	5-year
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Panel A Parameters Estimated from Regressions</i>										
Initial Relative Stock: $\hat{\rho}$	0.78 (0.13)	0.81 (0.12)	0.41 (0.12)	0.67 (0.14)	0.26 (0.09)	0.52 (0.12)	0.19 (0.06)	0.46 (0.17)	0.83 (0.03)	0.83 (0.03)
Initial Shifter: $\hat{\chi}$	0.84 (0.39)	1.11 (0.38)	1.00 (0.51)	1.07 (0.53)	0.63 (0.29)	1.04 (0.41)	0.38 (0.20)	1.07 (0.35)	1.66 (0.69)	2.06 (0.85)
Changes in Shifter: $\hat{\lambda}$	1.11 (0.54)	2.12 (0.67)	1.34 (0.51)	2.55 (0.90)	0.42 (0.29)	1.72 (0.63)	0.14 (0.16)	1.06 (0.65)	-1.58 (0.77)	-0.52 (0.95)
Spillovers:									-0.30 (0.15)	-0.23 (0.21)
Observations	232	125	345	165	345	165	149	78	3,459	3,447
Country covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Firm fixed effects									Yes	Yes
Industry \times Year fixed effects									Yes	Yes
Country \times Year fixed effects										Yes
<i>Panel B Implied Parameters</i>										
Long-run Effects	1.09	1.36	2.43	1.61	2.40	2.02	2.00	2.34	3.13	3.43
$\hat{\delta}$	0.22	0.19	0.59	0.33	0.74	0.48	0.81	0.54	0.47	0.40
$\hat{\sigma}$	1.68	1.88	1.41	1.70	1.22	1.53	1.15	1.47	1.78	1.85
$\hat{\varepsilon}$	4.06	4.96	2.86	4.13	2.01	3.37	1.66	3.12	4.50	4.81
$\hat{\delta}\hat{\sigma}$	0.38	0.35	0.83	0.57	0.91	0.74	0.93	0.80	0.84	0.74
<i>Panel C Equilibrium and Welfare Comparison (Baseline: $\bar{\tau} = 0.07$)</i>										
n^{SP}/n^{EQ}	0.83	0.82	0.56	0.76	0.40	0.66	0.34	0.60	0.47	0.61
$U^{SP} - U^{EQ}$	0.01	0.01	0.03	0.01	0.04	0.02	0.05	0.02	0.03	0.02
<i>Panel D Equilibrium and Welfare Comparison (Alternative: $\bar{\tau} = 0.03$)</i>										
n^{SP}/n^{EQ}	0.91	0.90	0.75	0.87	0.64	0.82	0.58	0.78	0.69	0.78
$U^{SP} - U^{EQ}$	0.002	0.002	0.01	0.003	0.01	0.004	0.01	0.01	0.01	0.01

Notes: This table presents regression estimates (Panel A), implied parameter values (Panel B) and implied distortions and welfare results (Panels C and D) for the automation application. Regressions are estimated with ordinary least squares and heteroscedasticity-robust standard errors clustered at country-level are presented in parentheses. All regressions are weighted by manufacturing employment in 1990. The dependent variable is relative number of newly granted patents for automation technologies relative to other utility patents divided by relative stock of patents related to automation relative to other utility patents (in logs, unless otherwise indicated). Shifters are expected 20-year level and change of the ratio of workers above the age of 56 to workers between 21 and 55 (in logs). Country covariates, included in columns 1-4, are region dummies, and the 1990 values of log GDP per capita, log of population, average years of schooling and the ratio of workers above 56 to workers aged 21 in 1990 interacted with period dummies. Columns 1 and 2 replicate the specifications from Table 2. Columns 3 and 4 use $\ln(1+x)$, while columns 5 and 6 use the inverse hyperbolic sine transformation. Columns 7 and 8 are for the OECD sample (with $\ln x$ as in our main specifications). Columns 9 and 10 report estimates from Dechezleprêtre et al.'s (2022) firm-level data, using a sample of firms with at least four automation patents. These regressions also include spillovers from country-level relative stock of knowledge. Column 9 controls for firm fixed effects and industry by year fixed effects, while column 10 additionally includes country by time fixed effects. The parameters δ and σ in these two columns are computed using the equations with spillovers provided above. Panel C uses 15% quasi-rents for workers, and Panel D uses 7.5% quasi-rents.

Table B2: Robustness for Health Application

LHS	$\ln(x)$	$\ln(x)$	$\ln(1+x)$	$\ln(1+x)$	$\operatorname{asinh}(x)$	$\operatorname{asinh}(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$
Samples	Full	Full	Full	Full	Full	Full	Full	Full	Full	Full	US	US
Frequency	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<i>Panel A Parameters Estimated from Regressions</i>												
Initial Relative Stock: $\hat{\rho}$	0.93 (0.03)	1.11 (0.03)	0.36 (0.02)	0.83 (0.04)	0.38 (0.02)	0.84 (0.03)	0.45 (0.04)	0.48 (0.04)	0.93 (0.03)	1.09 (0.03)	0.94 (0.10)	1.27 (0.12)
Initial Shifter: $\hat{\chi}$	0.10 (0.01)	0.14 (0.01)	0.05 (0.00)	0.11 (0.01)	0.07 (0.00)	0.13 (0.01)	0.07 (0.01)	0.10 (0.01)	0.11 (0.01)	0.14 (0.01)	0.32 (0.12)	0.26 (0.10)
Changes in Shifter: $\hat{\lambda}$	-0.004 (0.02)	0.001 (0.02)	-0.01 (0.01)	0.002 (0.01)	-0.01 (0.01)	0.002 (0.02)	-0.04 (0.03)	-0.06 (0.03)	-0.004 (0.02)	0.01 (0.02)	0.26 (0.17)	0.12 (0.07)
Observations	55,699	37,389	75,399	44,569	75,399	44,569	55,702	37,394	55,625	37,358	1,243	741
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes						
Disease fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes			Yes	Yes
Period-Country fixed effects									Yes	Yes		
Period-Disease fixed effects									Yes	Yes		
<i>Panel B Implied Parameters</i>												
Long-run Effects	0.11	0.14	0.15	0.13	0.18	0.16	0.16	0.22	0.11	0.14	0.34	0.26
$\hat{\delta}$	0.07	0.00	0.64	0.17	0.62	0.16	0.55	0.52	0.07	0.00	0.06	0.00
$\hat{\sigma}$	1.10	1.14	1.05	1.11	1.06	1.13	1.07	1.09	1.11	1.14	1.31	1.26
$\hat{\varepsilon}$	1.18	1.26	1.09	1.19	1.11	1.24	1.12	1.17	1.19	1.26	1.57	1.47
$\hat{\delta}\hat{\sigma}$	0.08	0.00	0.67	0.19	0.66	0.18	0.58	0.57	0.08	0.00	0.08	0.00
<i>Panel C Equilibrium and Welfare Comparison (Markups Only)</i>												
n^{SP}/n^{EQ}	0.43	0.45	0.11	0.39	0.11	0.38	0.17	0.17	0.43	0.45	0.37	0.42
$U^{SP} - U^{EQ}$	0.06	0.05	0.15	0.06	0.15	0.07	0.12	0.12	0.06	0.05	0.06	0.06
<i>Panel D Equilibrium and Welfare Comparison (Externalities Only)</i>												
n^{SP}/n^{EQ}	0.58	0.59	0.24	0.54	0.24	0.53	0.31	0.32	0.58	0.59	0.52	0.56
$U^{SP} - U^{EQ}$	0.18	0.17	0.48	0.21	0.47	0.21	0.39	0.38	0.18	0.17	0.22	0.19

Notes: This table presents regression estimates (Panel A), implied parameter values (Panel B) and implied distortions and welfare results (Panels C and D) for the health application. Regressions are unweighted and estimated with ordinary least squares and heteroscedasticity-robust standard errors clustered at country-level are presented in parentheses. Observations are at the country-disease-period level. The dependent variable is relative number of new medical articles for each disease divided by relative stock of medical articles for that disease (in logs, unless otherwise indicated). Columns 1 and 2 replicate the main specifications from Table 2. Columns 3 and 4 use $\ln(1+x)$, while columns 5 and 6 use the inverse hyperbolic sine transformation. Columns 7 and 8 drop the country fixed effects, while columns 9 and 10 include period times country and period times disease fixed effects. Columns 11 and 12 focus on just the US observations. Panel C considers the implications of markup differences, and Panel D depicts the implications of an externality estimate based on the shortfall of quality-adjusted life year gains from curative vs. preventative technologies.

Table B3: Robustness for Energy Application

LHS	$\ln(1+x)$	$\ln(1+x)$	$\operatorname{asinh}(x)$	$\operatorname{asinh}(x)$	$\ln(1+x)$	$\ln(1+x)$	$\ln(1+x)$	$\ln(1+x)$	$\ln(1+x)$	$\ln(1+x)$
Frequency	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Panel A Parameters Estimated from Regressions</i>										
Initial Relative Stock: $\hat{\rho}$	0.81 (0.03)	0.86 (0.04)	0.82 (0.02)	0.86 (0.04)	0.58 (0.03)	0.51 (0.04)	0.81 (0.03)	0.86 (0.05)	0.81 (0.03)	0.86 (0.04)
Initial Shifter: $\hat{\chi}$	-1.52 (0.29)	-1.06 (0.66)	-1.99 (0.36)	-1.44 (0.81)	-0.07 (0.09)	-0.48 (0.14)	-1.51 (0.28)	-2.02 (0.40)	-1.50 (0.29)	-1.14 (0.67)
Changes in Shifter: $\hat{\lambda}$	-0.45 (0.20)	1.12 (0.82)	-0.61 (0.26)	1.38 (1.00)	0.21 (0.13)	1.16 (0.36)	-0.47 (0.14)	-0.28 (0.21)	-0.43 (0.21)	1.09 (0.83)
Spillovers:									0.03 (0.03)	-0.07 (0.05)
Observations	13,648	6,824	13,648	6,824	13,648	6,824	13,648	6,824	13,648	6,824
Firm covariates	Yes	Yes	Yes	Yes	Yes	Yes			Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes			Yes	Yes	Yes	Yes
<i>Panel B Implied Parameters</i>										
Long-run Effects	-1.89	-1.23	-2.43	-1.68	-0.12	-0.94	-1.87	-2.34	-1.86	-1.32
$\hat{\delta}$	0.19	0.14	0.18	0.14	0.42	0.49	0.19	0.14	0.19	0.14
$\hat{\sigma}$	2.73	2.53	3.07	2.92	1.12	1.49	2.72	3.41	2.71	2.61
$\hat{\varepsilon}$	7.27	6.56	8.51	7.99	1.44	2.78	7.24	9.75	7.21	6.86
$\hat{\delta}\hat{\sigma}$	0.53	0.36	0.56	0.41	0.47	0.73	0.52	0.47	0.52	0.37
<i>Panel C Equilibrium and Welfare Comparison</i>										
n^{SP}/n^{EQ}	0.44	0.57	0.37	0.50	0.74	0.46	0.45	0.40	0.45	0.56
$U^{SP} - U^{EQ}$	0.03	0.02	0.04	0.03	0.01	0.03	0.03	0.04	0.03	0.02
<i>Panel D Equilibrium and Welfare Comparison (Using Global SCC)</i>										
n^{SP}/n^{EQ}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$U^{SP} - U^{EQ}$	13.74	8.94	16.99	11.60	3.50	12.15	13.62	15.58	13.55	9.50

Notes: This table presents regression estimates (Panel A), implied parameter values (Panel B) and implied distortions and welfare results (Panels C and D) for the health application. Regressions are unweighted and estimated with ordinary least squares and heteroscedasticity-robust standard errors clustered at firm level are presented in parentheses. Observations are at the firm-period level. The dependent variable is relative number of newly granted patents for dirty technologies relative to newly granted patents for clean technologies. Shifters are firm-level fuel prices adjusted (based on firm-level fuel consumption) inclusive of taxes (in $\ln(1+x)$ form unless otherwise indicated). All specifications include firm and period fixed effects as well as the values of government R&D subsidies for clean innovation, regulations over emissions, the relevant country's GDP per capita for that period (as in Aghion et al., 2016). Columns 1 and 2 replicate the main specifications from Table 2. Columns 3 and 4 use the inverse hyperbolic sine transformation. Columns 5 and 6 drop the firm fixed effects. Columns 7 and 8 additionally include the relative stock of knowledge in other firms in the same country as in the baseline specification of Aghion et al. (2016). In this case, we use the equations with spillovers for computing δ and σ . Panel C uses an externality number based on from Rennert et al.'s (2022) estimate of the social cost of CO2, converted to US-equivalent damages (see text for details), and Panel D uses their estimate for worldwide damages.

Appendix C: Data Sources and Construction

In this part of the Appendix, I describe the data sources for the three empirical exercises and provide some additional details.

Automation

Data for automation patents by country are directly from Acemoglu and Restrepo (2022). The flow of automation and non-automation patents were computed from the patents by different countries filed at the USPTO. In particular, all patents that are in the USPTO 901 class (technologies related to industrial robots) and all patents referencing this class are classified as automation patents. Aging variables are from the United Nations data, while the country-level covariates (GDP per capita, population, and average years of schooling) are from version 9.0 of the Penn World Tables (Feenstra, Inklaar and Timmer, 2015). Regressions are weighted by manufacturing value added in 1990 (sourced from the United Nations Industrial Development Organization).

Using these definitions, the exact estimating equations for columns 1 and 2 of Table 2 are:

$$\ln \left(\frac{\Delta n_{ct}}{n_{ct}} \right) = -\rho \ln n_{ct} + \chi \ln z_{ct} + \lambda \Delta \ln z_{ct} + \mathbf{X}_{c,1990} \boldsymbol{\gamma}_t + \epsilon_{ct}, \quad (\text{C1})$$

where Δn_{ct} is the ratio of the flow relative automation patents (compared to non-automation patents) and n_{ct} is the relative technology stock (automation patent stock relative to non-automation patents stock). Stocks are computed from the corresponding flow variables using a 20% depreciation rate, as explained in the text.

The forcing variables are: the (log of) the ratio workers aged 56 and above to those between the ages of 25 and 55 and the (log of) 15- or 20-year ahead change in the ratio workers age 56 and above to those between the ages of 25 and 55. Finally, $\mathbf{X}_{c,1990}$ denotes the country covariates (log GDP per capita, log population, and average years of schooling, all in 1990), and the fact that its coefficient is time varying designates that these covariates are allowed to have a separate effect in every time period. The sample covers 69 countries and the time period 1986-2015. The observations are weighted by value added in manufacturing in 1990 and standard errors are clustered by country. The models estimated in Table B1 are variations of these equations as explained in the text.

In addition, the estimates in columns 11 and 12 of Table B1 are provided directly by Dechezleprêtre et al. (2022), based on their firm-level data set on automation and non-automation patents. The reader is referred to their paper for variable definitions and sources.

Medical Research

Our estimates for medical research's responsiveness to disease burden, depicted in columns 3 and 4 of Table 2 and Table B2, come directly from Acemoglu, Moscona, Sastry and Williams (2023). The estimating equation is similar to (24) in the text:

$$\ln \left(\frac{\Delta N_{dct}}{N_{dct}} \right) = \eta_d + \Gamma_c + \Upsilon_t - \frac{\rho}{2} \ln N_{dct} + \chi \ln Z_{dct} + \lambda \Delta \ln Z_{dct} + \epsilon_{dct}, \quad (\text{C2})$$

where N_{dct} (ΔN_{dct}) is the stock (flow) of medical scientific articles on disease d in country c at time t . Stocks are again computed from flows using a 20% depreciation rate. The forcing variables are the level and change of disease burdens, defined as declines in the number of disability-adjusted life years caused

by a disease in a country and time period in our sample. These calculations are based on data from the Global Burden of Disease (GBD) project. Finally, η_d , Γ_c and Υ_t are, respectively, disease, country and time fixed effects, and in some specifications, two-way fixed effects are also included. All regressions in this case are unweighted. Additional details can be obtained from Acemoglu et al. (2023).

Energy

Data on the relationship between fuel prices and automobile patents come directly from Aghion et al. (2016). The data on flows of patents are based on the World Patent Statistical Database (PATSTAT) maintained by the European Patent Office (EPO), and innovation is measured using a count of patents by application/filing date. The authors use data on tax-inclusive fuel prices, from the International Energy Agency (IEA), to compute a time-varying, country-specific fuel price by averaging the prices of diesel and gasoline prices. Country-specific fuel prices are then used to construct firm-level fuel prices as a weighted average of fuel prices across countries based on the firm's expected market share across countries (in practice, using a time-invariant share of the firm's sales in each market). The covariates are log GDP per capita (sourced from World Development Indicators), log R&D subsidies (from the IEA), and exposure to air pollution regulations. Emission regulations are for maximum level of tailpipe emissions for pollutants for new automobiles, coded between 0 and 5, and are taken from Dechezleprêtre et al. (2012).

The exact estimating equation for columns 5 and 6 of Table 2 is similar to but a little different from (C1). In particular, Aghion et al. (2016) impute log patent stocks as zero when stocks are zero. We avoid this by using $\ln(1+x)$ consistently for both flow and stock variables throughout this application. This gives our estimating equation as:

$$\ln\left(\frac{\Delta\tilde{n}_{fct}}{\tilde{n}_{fct}}\right) = -\rho\ln(\tilde{n}_{fct}) + \chi\ln z_{fct} + \lambda\Delta\ln z_{fct} + \mathbf{X}_{fct}\boldsymbol{\gamma} + \epsilon_{ct}, \quad (\text{C3})$$

where

$$\ln\left(\frac{\Delta\tilde{n}_{fct}}{\tilde{n}_{fct}}\right) = \ln\left(\frac{1 + \text{Patent}_{fct}^{\text{clean}}}{1 + \text{Patent}_{fct}^{\text{dirty}}}\right) - \ln\left(\frac{1 + \text{Stock}_{fct}^{\text{clean}}}{1 + \text{Stock}_{fct}^{\text{dirty}}}\right)$$

and likewise,

$$\ln(\tilde{n}_{fct}) = \ln\left(\frac{1 + \text{Stock}_{fct}^{\text{clean}}}{1 + \text{Stock}_{fct}^{\text{dirty}}}\right),$$

with $\text{Patent}_{fct}^{\text{clean}}$ and $\text{Patent}_{fct}^{\text{dirty}}$, respectively, denoting the flow of clean and dirty automobile patents for firm f located in country c at time t , and $\text{Stock}_{fct}^{\text{clean}}$ and $\text{Stock}_{fct}^{\text{dirty}}$ likewise denoting the stocks of clean and dirty patents. The forcing variables, as described above, are based on firm-level fuel prices and their changes, while covariates are now time-varying but have constant coefficients. Regressions are unweighted and estimated by ordinary least squares, and standard errors are heteroscedasticity-robust and clustered at the country level. The models estimated in Table B3 are variations of these equations as explained in the text.

Markup Estimation

In this part of the Appendix, I describe our markup estimation strategies. Throughout, each firm is assumed to have a single, well-defined price at each point in time. Then, the gross markup of firm i at

time t is defined as

$$\Lambda_{it} = \frac{P_{it}}{MC_{it}}, \quad (\text{C4})$$

where P_{it} is this firm's price at time t and MC_{it} is its marginal cost. Note that in the text I focused on net markups defined as

$$\mu_{it} = \Lambda_{it} - 1.$$

Production Function Estimation Methods

The production function method follows De Loecker et al. (2020). Let us first focus on a single industry, and suppose that each firm i in this industry has a production function

$$Q_{it}(V_{it}, K_{it}) \quad (\text{C5})$$

at time t , with V_{it} denoting a composite of variable inputs (labor and material) and K_{it} representing its capital stock. Suppose that the capital stock is a quasi-fixed factor, meaning that it is chosen in advance (and hence the designation of the other factors as “variable”). The function Q_{it} is firm and time-varying, for example, it includes information on the firm's (revenue) productivity upon which variable costs may depend. In the estimation, the function Q_{it} will be taken to be Cobb-Douglas.

Consider the elasticity of this production function with respect to variable inputs, V_{it} , denoted by θ_{it}^V :

$$\theta_{it}^V = \frac{\partial Q_{it}}{\partial V_{it}} \frac{V_{it}}{Q_{it}} = \frac{1}{MC_{it}} \frac{P_{it}^V V_{it}}{Q_{it}},$$

where P_{it}^V is the price of the composite variable input, and the second equality exploits the fact that, because K_{it} is fixed, the marginal cost of production is $MC_{it} = P_{it}^V / (\partial Q_{it} / \partial V_{it})$. Next, using the definition of the markup in (C4) to substitute out MC_{it} and rearranging, we obtain

$$\Lambda_{it} = \theta_{it}^V \frac{P_{it} Q_{it}}{P_{it}^V V_{it}}.$$

Given this equation, firm-level markups can be estimated with data on revenue, $P_{it} Q_{it}$, cost of variable inputs, $P_{it}^V V_{it}$, and crucially the elasticity of the firm's production function with respect to variable inputs, θ_{it}^V .

Here I briefly outline De Loecker et al.'s (2020) estimation strategy, which I follow. Recall that the capital stock is quasi-fixed. Suppose also that observed sales are given by $\text{Sales}_{it} = \varepsilon_{it} Q_{it}(V_{it}, K_{it})$, where ε_{it} is a demand shifter realized after all input decisions are made. Finally, as noted below, suppose that the function Q_{it} is Cobb-Douglas, and denote the Hicks-neutral productivity of firm i at time t by Ω_{it} . Then we have

$$\ln \text{Sales}_{it} = \theta_t^V \ln V_{it} + \theta_t^K \ln K_{it} + \ln \Omega_{it} + \varepsilon_{it}, \quad (\text{C6})$$

which allows for the Cobb-Douglas exponents, and thus output elasticities, to be time-varying, but constant across firms (within the industry being considered). The difficulty in the estimation of (C6) is that the firm knows Ω_{it} when choosing its composite variable input V_{it} , and thus OLS estimation will lead to biased output elasticities. De Loecker et al. (2020) deal with this problem by using a control function approach based on Olley and Pakes (1996). For example, Hicks-neutral productivity Ω_{it} can be assumed to be measurable with respect to the firm's capital stock K_{it} , investment I_{it} , and additional control variables

related to factor demands denoted by Z_{it} . This implies a relationship of the form

$$\ln(\Omega_{it}) = \phi_t(\ln K_{it}, \ln I_{it}, Z_{it}),$$

so that the elasticity of output with respect to variable inputs, θ_t^V , can be estimated from the following equation:

$$\ln \text{Sales}_{it} = \theta_t^V \ln V_{it} + \theta_t^K \ln K_{it} + \phi_t(\ln K_{it}, \ln I_{it}, Z_{it}) + \varepsilon_{it}. \quad (\text{C7})$$

I follow De Loecker et al. (2020) and include the following terms in the ϕ function: a quadratic and cubic in $\ln K_{it}$, a main, quadratic and cubic in $\ln I_{it}$, and the interaction between these two variables, $\ln K_{it} \ln I_{it}$.²⁸ In addition, as in their specification, the Z_{it} variable includes the ratio of the firm's total costs to the four-digit industry total cost, and the ratio of the firm's total costs to the economy-wide total cost.

Once estimates of the variable input elasticity $\hat{\theta}_t^V$ are obtained, (gross) markups can be computed as

$$\hat{\Lambda}_{it}^P = \hat{\theta}_t^V \frac{P_{it} Q_{it}}{P_{it}^V V_{it}}, \quad (\text{C8})$$

where the superscript P specifies that this is a markup estimated using the production function method. The (net) markup is then $\hat{\mu}_{it}^P = \hat{\Lambda}_{it}^P - 1$.

De Loecker et al.'s (2020) baseline estimates are based on a variant based on Akerberg, Caves and Frazer (2015), where the composite variable input is used instead of investment and $\ln(\Omega_{it})$ is assumed to follow a first-order Markov process. For this specification, I directly use their estimates of these elasticities, reported in De Loecker et al. (2020a).

One drawback of the production function method is that the estimation of θ_{it}^V requires the model and the measurability assumptions embedded in the control function to be correctly specified.

As an alternative, De Loecker et al. (2020) also use cost shares to estimate θ_t^V . In particular, they compute industry-level output elasticities as

$$\tilde{\theta}_t^V = \text{median} \left\{ \frac{P_{it}^V V_{it}}{P_{it}^V V_{it} + R_t K_{it}} \right\},$$

where the median is across all firms within a two-digit industry, and R_t is the user cost of capital. In this case, (gross) markups can be obtained as

$$\hat{\Lambda}_{it}^C = \tilde{\theta}_t^V \frac{P_{it} Q_{it}}{P_{it}^V V_{it}},$$

where the superscript C refers to the fact that output elasticities are now estimated from cost shares (but still taken to be common across firms within an industry). Then, naturally, $\hat{\mu}_{it}^C = \hat{\Lambda}_{it}^C - 1$. In practice, this approach gives similar results to the production function estimation.

One drawback in this case is that, although the functional form assumptions of the production function method are relaxed, the assumption that there is a common θ_t^V at the industry level is challenged by the fact that there is a large variation in cost shares, and the median is an arbitrary way of resolving this issue.

Another drawback of both approaches from my point of view is that it is not entirely clear whether markups relative to variable costs is the right notion for μ in the model, since this parameter captures

²⁸As in their paper, investment in Compustat is computed from the capital stock data assuming a 10% depreciation rate, that is, $I_{it} = K_{it} - 0.9 \cdot K_{it-1}$.

how profitable a technology is and regulates incentives for innovation and entry. If one technology is more capital intensive and has higher markup estimates using the production function methods used by De Loecker et al. (2020), it may nonetheless have lower profitability and lower μ in the sense of the theoretical framework in the text. This motivates the next strategy.

The Accounting Method

As an alternative, one could directly use an estimate of profits to compute markups.

Using the same notation and terminology, (gross) markups are estimated in this case as

$$\hat{\Lambda}_{it}^A = \frac{P_{it}Q_{it}}{P_{it}^V V_{it} + R_t K_{it}}, \quad (\text{C9})$$

where R_t is the user cost of capital (again assumed to be the same across firms in the industry). We then have: $\hat{\mu}_{it}^A = \hat{\Lambda}_{it}^A - 1$.

The drawback of this approach is well known: accounting profits do not correspond to economic profits. The advantage, on the other hand, is related to the discussion at the end of the previous subsection. This method takes into account capital costs explicitly, and thus may be more informative about the overall profitability of a technology/subsector.

Data

I follow De Loecker et al. (2020) and use Compustat North America for firm-level markup estimation. Compustat Fundamentals Annual extract is obtained through Wharton Research Data Service (WRDS), and I use the same variables as De Loecker et al. (2020a). Namely, the variable *SALE* measures revenues and variable costs are measured using the variable *COGS* (cost of goods sold, which includes expenses for materials, labor, overhead and other intermediate inputs). The capital stock of each firm is measured using the variable *PPEGT* (property, plant, and equipment gross total). The user cost of capital is also computed as in their paper: $R_t = \text{nominal interest rate}_t - \text{inflation}_t + \text{depreciation rate}$.^{29, 30, 31} I set the depreciation rate at 10%. We exclude firms in the top and bottom 1% of cost of goods to sales ratio (*COGS/SALE*) and cost-shares, which are likely to have extreme values due to measurement error.³²

Aggregating Markups

Throughout, I aggregate firm-level markups to industry-group level (in this instance, preventative health care vs. curative health care) by using the ratio of firm costs to industry-group costs. Specifically, our main estimates aggregate (gross) markups with the following equation:

$$\Lambda_{Jt}^P = \sum_{i \in J} \frac{\text{Cost}_{it}}{\text{Cost}_{jt}} \times \Lambda_{it}^P,$$

²⁹I follow De Loecker et al. (2020) and use the federal funds rate, FEDFUNDS, and the annual percent change in the relative price of investment goods, PIRIC. Both variables are taken from the Federal Reserve Economic Data, FRED.

³⁰Board of Governors of the Federal Reserve System (US), Federal Funds Effective Rate [FEDFUNDS], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/FEDFUNDS>, March 26, 2023.

³¹Relative Price of Investment Goods [PIRIC], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PIRIC>, March 27, 2023.

³²In particular, for this exercise, cost shares are measured as $\frac{COGS}{COGS+KEXP}$ and $\frac{COGS}{COGS+KEXP+SGA}$, where *SGA* measures selling, general, and administrative expenses

where $\text{Cost}_{it} = P_{it}^V V_{it} + R_t K_{it}$, and J denotes the industry-group in question (with $i \in J$ designating that firm i belongs to this group) such that $\text{Cost}_{jt} = P_{jt}^V V_{jt} + R_t K_{jt}$. I choose cost-based aggregation rather than using revenue-weights as in De Loecker et al. (2020), since, as these authors also recognize, revenue-based estimation can lead to inflated aggregate markups because high markup firms, which generate higher revenues, receive greater weights.

Firm Classification

This subsection explains how firms in Compustat are assigned to preventative and curative health care.

The classification is on the basis of the main North American Standard Industry Code (NAICS) assigned to firms in Compustat.³³

Preventative: Health care firms whose main activity is in basic health provision, diagnosis or manufacture of vaccines and related products are assigned to the preventative health care group.³⁴ Firms with the following main NAICS codes are included in this category:

- *NAICS 621 - Ambulatory health services:* Firms that provide health care services directly or indirectly to ambulatory patients and do not usually provide inpatient services. Includes outpatient services provided by physicians, dentists, and other health practitioners. Also includes outpatient care centers, medical and diagnostic laboratories, home health care services, and other ambulatory health care services.
- *NAICS 325413 - In-vitro diagnostic substances manufacturing:* Firms that manufacture in-vitro (i.e., not taken internally) diagnostic substances (chemical, biological, or radioactive substances). Substances are used for diagnostic tests, such as blood glucose, HIV, pregnancy, and other tests. It also involves manufacturing hematology, hormone, microbiology, and viral diagnostic substances, among others.
- *NAICS 325414 - Biological product (except diagnostic) manufacturing:* Firms primarily involved in manufacturing vaccines, toxoids, blood fractions, etc.

Curative: Health care firms whose main activity is in pharmaceutical preparation and high-tech medical equipment manufacturing firms (including a few that are related to advanced diagnostics) are assigned to the curative health care group. These firms are again identified based on their main NAICS codes, including the following categories:

- *325412 - Pharmaceutical preparation manufacturing:* Firms manufacturing in-vivo diagnostic substances and pharmaceutical preparations (except biological) intended for internal and external consumption in dose forms, such as tablets, capsules, vials, ointments, powders, solutions, and suspensions.
- *334510 - Electromedical and electrotherapeutic apparatus manufacturing:* Firms manufacturing electromedical and electrotherapeutic apparatus such as magnetic resonance imaging equipment, medical ultrasound equipment, pacemakers, hearing aids, electrocardiographs, and electromedical endoscopic equipment.

³³Codes and descriptions obtained from the U.S. Census Bureau North American Industry Classification System (NAICS) at <https://www.census.gov/naics/?99967>.

³⁴In addition, preventative health care should also include those in the area of social assistance, *NAICS 624*, which comprises firms providing individual and family services, community food and housing, vocational rehabilitation services, child daycare services, as well as emergency and other relief services. Nevertheless, there are no firms in this NAICS category in Compustat.

- *339112 - Surgical and medical instrument manufacturing*: Firms manufacturing medical, surgical, ophthalmic, and veterinary instruments and apparatus (except electrotherapeutic, electromedical, and irradiation apparatus). Examples are syringes, needles, anesthesia apparatus, blood transfusion equipment, catheters, surgical clamps, and medical thermometers.
- *339113 - Surgical appliance and supplies manufacturing*: Firms manufacturing surgical appliances and supplies such as orthopedic devices, prosthetic appliances, surgical dressings, personal safety equipment, hospital beds, operating tables, etc.


Table [D1](#), included in Appendix [D](#), which is available upon request, provides a full list of health care firms in Compustat and their assignment to preventative and curative categories. It also lists the relevant sample period for the firm, sales, costs of goods sold, and capital stock, as well as the four measures of (net) markup—based on production function, cost share and accounting methods. In total, our sample includes 658 preventative and 1,069 curative health care firms. At the bottom of each panel, (cost-weighted) averages of the markups are also presented.

The four panels of Figure [C1](#) show the evolution of markups based on the two production function estimation methods, cost share and accounting methods, separately for firms in the preventative and curative categories. Each panel also gives the average (net) markup, which corresponds to μ in the model. The trends are fairly similar with the different methods and show some fluctuations and also a significant increase in markups among firms in the curative category. This is consistent with the patterns reported by De Loecker et al. (2020) at a higher level of aggregation. The increase in markups among curative firms is in fact larger than at the two-digit level patterns depicted by De Loecker et al. (2020).


Our baseline uses the averages in Panel a, which give $\mu_1 = 1.70$ for preventative firms and $\mu_2 = 0.46$ for curative firms as also shown in Table [1](#) in the text. The numbers in the other panels are quite similar, and using these numbers instead yields broadly similar results to those reported in Table [2](#). Table [C1](#) shows the technology ratio and welfare loss estimates corresponding to the specifications in Table [B2](#), if instead we use the markup estimates in Panels b, c or d of Figure [C1](#).

Figure C1: Aggregate Markups by Sector Group


a. Total cost-weighted average of firm-level μ_{it}^{P1}
(first production function estimation method)




b. Total cost-weighted average of firm-level μ_{it}^{P2}
(second production function estimation method)



c. Total cost-weighted average of firm-level μ_{it}^C
(cost-share estimation method)



d. Total cost-weighted average of firm-level μ_{it}^A
(accounting method)



Note: This figure depicts total cost-weighted averages of firm-level markups across the preventative and curative technology groups. Cost shares are defined as $(COGS_{it} + R_t PPEG T_{it}) / (COGS_{jt} + R_t PPEG T_{jt})$ (see Appendix D Table D1). The four panels use firm-level markups μ_{it}^{P1} , μ_{it}^{P2} , μ_{it}^C and μ_{it}^A , which are based, respectively, on the first and second production function estimation methods, cost-share estimation method and the accounting method, as described in Section 6. The list of firms is given in Appendix D Table D1.

Table C1: Sensitivity Analysis of Technology Distortions and Welfare Losses from Markups

LHS	$\ln(x)$	$\ln(x)$	$\ln(1+x)$	$\ln(1+x)$	$\operatorname{asinh}(x)$	$\operatorname{asinh}(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$	$\ln(x)$
Samples	Full	Full	Full	Full	Full	Full	Full	Full	Full	Full	US	US
Frequency	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year	5-year	10-year
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<i>Panel A Parameters Estimated from Regressions</i>												
Initial Relative Stock: $\hat{\rho}$	0.93 (0.03)	1.11 (0.03)	0.36 (0.02)	0.83 (0.04)	0.38 (0.02)	0.84 (0.03)	0.45 (0.04)	0.48 (0.04)	0.93 (0.03)	1.09 (0.03)	0.94 (0.10)	1.27 (0.12)
Initial Shifter: $\hat{\chi}$	0.10 (0.01)	0.14 (0.01)	0.05 (0.00)	0.11 (0.01)	0.07 (0.00)	0.13 (0.01)	0.07 (0.01)	0.10 (0.01)	0.11 (0.01)	0.14 (0.01)	0.32 (0.12)	0.26 (0.10)
Changes in Shifter: $\hat{\lambda}$	-0.004 (0.02)	0.001 (0.02)	-0.01 (0.01)	0.002 (0.01)	-0.01 (0.01)	0.002 (0.02)	-0.04 (0.03)	-0.06 (0.03)	-0.004 (0.02)	0.01 (0.02)	0.26 (0.17)	0.12 (0.07)
Observations	55,699	37,389	75,399	44,569	75,399	44,569	55,702	37,394	55,625	37,358	1,243	741
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes						
Disease fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes			Yes	Yes
Period-Country fixed effects									Yes	Yes		
Period-Disease fixed effects									Yes	Yes		
<i>Panel B Implied Parameters</i>												
Long-run Effects	0.11	0.14	0.15	0.13	0.18	0.16	0.16	0.22	0.11	0.14	0.34	0.26
$\hat{\delta}$	0.07	0.00	0.64	0.17	0.62	0.16	0.55	0.52	0.07	0.00	0.06	0.00
$\hat{\sigma}$	1.10	1.14	1.05	1.11	1.06	1.13	1.07	1.09	1.11	1.14	1.31	1.26
$\hat{\varepsilon}$	1.18	1.26	1.09	1.19	1.11	1.24	1.12	1.17	1.19	1.26	1.57	1.47
$\hat{\delta}\hat{\sigma}$	0.08	0.00	0.67	0.19	0.66	0.18	0.58	0.57	0.08	0.00	0.08	0.00
<i>Panel C Equilibrium and Welfare Comparison (Markups Only), Baseline</i>												
n^{SP}/n^{EQ}	0.43	0.45	0.11	0.39	0.11	0.38	0.17	0.17	0.43	0.45	0.37	0.42
$U^{SP} - U^{EQ}$	0.06	0.05	0.15	0.06	0.15	0.07	0.12	0.12	0.06	0.05	0.06	0.06
<i>Panel D Equilibrium and Welfare Comparison (Markups Only), Production function estimation method 2</i>												
n^{SP}/n^{EQ}	0.44	0.46	0.12	0.40	0.12	0.39	0.18	0.18	0.44	0.46	0.38	0.43
$U^{SP} - U^{EQ}$	0.05	0.05	0.14	0.06	0.14	0.06	0.11	0.11	0.05	0.05	0.06	0.05
<i>Panel E Equilibrium and Welfare Comparison (Markups Only), Cost-share based estimation</i>												
n^{SP}/n^{EQ}	0.50	0.51	0.16	0.45	0.16	0.45	0.23	0.23	0.50	0.51	0.44	0.48
$U^{SP} - U^{EQ}$	0.04	0.04	0.10	0.04	0.10	0.05	0.08	0.08	0.04	0.04	0.05	0.04
<i>Panel F Equilibrium and Welfare Comparison (Markups Only), Accounting method</i>												
n^{SP}/n^{EQ}	0.53	0.55	0.19	0.49	0.20	0.48	0.26	0.27	0.53	0.55	0.47	0.52
$U^{SP} - U^{EQ}$	0.03	0.03	0.08	0.04	0.08	0.04	0.07	0.07	0.03	0.03	0.04	0.03

Notes: This table shows how technology distortions and welfare losses change across the specifications considered in Table B2 for different values of markups. Panels A and B replicate the same panels in Table B2. The remaining four panels correspond to the four sets of markup estimates, $\hat{\mu}_{it}^{P1}$, $\hat{\mu}_{it}^{P2}$, $\hat{\mu}_{it}^C$ and $\hat{\mu}_{it}^A$, which are, respectively, from the first and second production function estimation methods, cost-share estimation method and the accounting method. See Figure C1 and Section 6 on the markup estimates, and see Table B2 on the parameter estimates and the underlying regression models for the different columns here.

Quality-Adjusted Life Years

In this section, we describe how differences between preventative and curative technologies in terms of quality-adjusted life years (QALYs) are estimated.

Methodology

Quality-Adjusted Life Years (QALYs) are a common measure used for evaluating the effectiveness of medical treatments and interventions. They quantify the overall gains in quantity and quality of life. QALYs are calculated by multiplying the number of years of life gained by a quality of life scale, which ranges from 0 (death) to 1 (perfect health). To access cost-effectiveness analyses in a comprehensive manner, we use the Cost-Effectiveness Analysis (CEA) Registry by the Center for the Evaluation of Value and Risk in Health, Tufts Medical Center. This registry includes studies on a wide range of health interventions, including drugs, medical devices, diagnostic tests, and prevention strategies and reports detailed information on the methods used in and results of each study.

We restrict the sample to modern healthcare innovations with studies conducted in the United States and benchmark the relevant innovation to the year of Food and Drug Administration (FDA) approval. We exclude a large number of studies included in the registry that evaluate the effectiveness of immunization drives and information campaigns. We focus on studies on pharmaceuticals, medical devices, and surgical procedures, especially those that compare a drug to placebo or no treatment. For these innovations we extract the QALYs gained per patient from the relevant journal article or website containing the study. In the case that a drug is compared to another drug instead of a placebo or no treatment, where possible, we search for auxiliary studies that compare one of the drugs in the main study to placebo and use that as a reference point to impute the effect and cost of all drugs in the main study relative to placebo.³⁵ Note that the QALY numbers obtained from this procedure can be negative, if new procedures are worse than no treatment or placebo, and are indeed so in a few cases.

Most estimates give QALY gains per patient. To construct comparable social benefits, I convert these estimates into QALY gains per dollar. Specifically, I use the following equation for each innovation i :

$$\text{QALY per dollar}_i = \frac{\text{QALY per patient}_i \times \text{Number of users}_i}{\text{Cost per user}_i \times \text{Number of users}_i + \text{R\&D costs}_i}. \quad (\text{C10})$$

Intuitively, this expression corresponds to total benefits divided by total costs, including R&D costs. In estimating the number of users, I limit the horizon for each innovation to 20 years, which amounts to assuming that this innovation will be replaced by a new one on average every 20 years. Given these estimates, I construct the average quality-adjusted life year gains by preventative and curative technology groups as

$$\text{QALY per dollar}_G = \sum_{i \in G} \text{Cost share}_i \times \text{QALY per dollar}_i,$$

where G is either the preventative or the curative technology group, and

$$\text{Cost share}_i = \frac{\text{Cost per user}_i \times \text{Number of users}_i + \text{R\&D costs}_i}{\sum_{i' \in G} (\text{Cost per user}_{i'} \times \text{Number of users}_{i'} + \text{R\&D costs}_{i'})}.$$

In these equations, R&D costs are estimated from the medical literature, which provides average of

³⁵In principle, one might wish to obtain QALYs relative to a single dominant treatment that exists before the innovation. In practice, this did not prove to be straightforward, and hence I opted for making all comparisons relative to placebo or no treatment.

R&D costs by class of drugs, e.g., oncological, immunomodulant, therapeutic recombinant proteins and mAbs, cardiovascular, etc. The medical papers in this literature use a variety of methods to obtain R&D costs, including using proprietary databases with cost information at the individual drug level, mandatory SEC filings, and industry surveys. Virtually all papers involve accounting for both failed and approved drugs, the type and duration of clinical trials, and the status of drug review at the FDA such as fast track, accelerated approval, or priority review. These papers are helpfully reviewed in Table 1 of Schlander et al. (2021). We match each innovation to its pharmaceutical category and impute the cost of R&D as the average cost for that group of drugs. For example, amlodipine is a calcium channel blocker that can treat high blood pressure and chest pain. As it acts on the cardiovascular system, we impute its R&D cost as the average R&D cost for all cardiovascular drugs. For surgical procedures, we use the sum of R&D expenses over several years or total invested capital for the primary manufacturer of equipment used in the surgical procedure, whichever data are available.

Per-patient usage costs are taken from the same papers that present the QALY benefits. These costs are often constructed as a sum of a direct treatment cost and an indirect health care cost, which imputes a production loss due to the patient’s injury and inability to work. Note that both the QALY estimates and per-patient usage cost are relative to placebo or no treatment, and thus we obtain negative values in a few cases. This is primarily because treatment avoids other costs patients incur in the future.

We use three methodologies to estimate the number of users. First, we look for direct estimates of the number of patients using the drug. Such statistics are available on clinical.com or the Centers for Disease Control and Prevention (CDC) website. Second, if no direct estimate is available, we estimate the number of users by dividing the US total sales for the innovation by the annual therapy cost, which is itself the product of dosage, frequency, and price per dosage. Lastly, if the direct and revenue imputation approach are infeasible, we gauge the number of patients by multiplying the incidence of disease by the proportion of patients who undergo treatment by the innovation.

For one-time innovations, such as some surgical procedures, we focus on the number of annual patients and multiply this by 20. For innovations that involve recurring use, such as antihypertensive medication, the number of patients is given by the contemporaneous usage prevalence, under the assumption that a patient uses the drug for the duration of the time horizon.

Finally, given the QALY per dollar_G estimates, we set $\tilde{\tau}_1 = 0$, and compute $\tilde{\tau}_2$ on the basis of the relative shortfall of the curative technologies compared to the preventative technologies:

$$\tilde{\tau}_2 = 1 - \frac{\text{QALY per dollar}_{\text{curative}}}{\text{QALY per dollar}_{\text{preventative}}}.$$

List of Procedures

I now provide further details on the procedures and innovations selected in the computation of the QALY numbers.

The CEA Registry contains roughly 10,000 entries for the United States with QALY outcomes. First, we manually considered each study to determine if it studied an innovation which constituted a modern healthcare innovation. This step eliminated immunization drives and information campaigns. Second, we excluded studies without a placebo or no treatment comparison. This step dropped a significant proportion of the sample, including studies comparing dosages of a given drug, evaluating different drugs for a given disease, or assessing the most efficacious treatment combination of a set of drugs (many of these studies are aimed at better informing *clinical* use, which is very different from our purpose here). Third, as screening procedures are not easily categorized into preventative vs. curative, we dropped all

such studies. Lastly, we sought the background and history of each innovation and kept those which were commercially developed in the late 20th century.

We then performed a second pass where we actively searched for cost-effectiveness studies relating to the top 20 drugs in the United States, as listed in clinicaltrials.gov. While data were not available for all 20 drugs, we were able to add 9 additional important drugs to our list.

Table D2 in Appendix D (available upon request) lists the 71 procedures we consider. In each case we provide reference to the source article where the medical information is taken, and list our estimates of R&D costs, usage costs, total QALY benefits, and our final QALY benefits per dollar. The two panels correspond to curative and preventative technologies.

Additional References for Appendix C

Akerberg, Daniel, Kevin Caves and Garrett Frazer (2015) “Identification Properties of Recent Production Function Estimators.” *Econometrica*, 83(6): 2411-2451.

Dechezleprêtre, Antoine, Richard Perkins, and Eric Neumayer (2012) “Regulatory Distance and the Transfer of New Environmentally Sound Technologies: Evidence from the Automobile Sector.” Working Paper no. 2012.33, Fondazione Eni Enrico Mattei, Milan.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020a) “Replication Data for: The Rise of Market Power and the Macroeconomic Implications.” *The Quarterly Journal of Economics*, Harvard Dataverse, V1, <https://doi.org/10.7910/DVN/5GH8XO>.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (2015) “The Next Generation of the Penn World Table.” *American Economic Review*, 105(10): 3150-82.

Olley G. Steven and Ariel Pakes (1996) “The Dynamics of Productivity in the Telecommunication Equipment Industry.” *Econometrica*, 64(6): 1263-1297.

Schlander, M., Hernandez-Villafuerte, K., Cheng, C. Y., Mestre-Ferrandiz, J., & Baumann, M. (2021) “How Much Does It Cost to Research and Develop a New Drug? A Systematic Review and Assessment.” *Pharmacoeconomics*, 39, 1243-1269.