



A Note on Teaching Automation in Undergraduate Courses

A growing macroeconomic literature argues that some new technologies are close substitutes for human labor and that the introduction of these technologies drives down wages and employment. This is at odds with the standard model taught to undergraduates, where technological progress and capital accumulation increase the marginal product of labor and the demand for workers, leading to higher wages and employment. I argue that insights from this new academic literature can easily be incorporated into the standard supply and demand framework and used to improve the discussion of labor markets in undergraduate courses.

Gregory Casey[†]

[†]Williams College

1. Introduction

There is a large gap between the macroeconomics we teach undergraduates and the macroeconomics that guides research and policy discussions. For example, undergraduate education on economic growth focuses on capital accumulation the Solow (1956) model, while research on growth focuses on other factors, like endogenous productivity and institutions (Jones & Romer, 2010). Undergraduate education on the conduct of monetary policy often focuses on the supply and demand for money, even though so-called “unconventional” monetary policies, like interest on reserves, are now the standard tools of the Federal Reserve (Ihrig & Wolla, 2022). A growing literature attempts to update undergraduate macroeconomics education to more accurately reflect the views of academic economists and the current policy environment (e.g., Acemoglu, 2013; Solis-Garcia, 2018; Hoyt & McGoldrick, 2019; Ihrig & Wolla, 2023). This paper contributes to this effort by revisiting how the labor markets are covered in undergraduate macroeconomics courses. The tools discussed in this paper are aimed at intermediate macroeconomics courses without calculus, but could also be used in advanced principles or labor economics courses.¹

How do new technologies affect wages and employment? This question has long garnered significant attention in public debates.² Unfortunately, this topic cannot be adequately addressed with the standard approach to teaching labor markets in undergraduate macroeconomics courses. In the standard approach, new technologies are assumed to increase the productivity of workers, shifting out the labor demand curve and leading to higher wages and employment. This is certainly a potential outcome of technological advance, but it is not the only one.

Recent advances in the macroeconomics literature argue that some new technologies decrease wages and employment for some workers (e.g., Acemoglu & Restrepo, 2018a, 2019, 2020). These advances in the literature include both empirical evidence and new models. The models adopt a task-based framework (e.g., Zeira, 1998; Acemoglu & Autor, 2011), where technological advances create new modes of production in which non-human inputs to production (e.g., “robots” or “artificial intelligence”³) can perfectly substitute for human labor, leading to lower wages and employment. I will refer to these technologies broadly as “automation technologies.” The gap between cutting-edge research and undergraduate instruction implies that students are being given an incomplete view of how economists think about labor markets. The gap also implies that students will not be able to use their economic training to engage meaningfully in public discussions about automation.

In this paper, I discuss how a small modification to the standard supply and demand model allows students to think about the implications of automation technologies. To capture the theoretical notion that some new technologies are perfect substitutes for human labor, the standard model is augmented to include a joint supply curve for humans and “robots”. Characterizing the equilibrium requires examining the intersection of the demand curve with two different supply curves, the human supply curve and the joint supply curve. The introduction of robots decreases wages and human employment. This outcome is not possible with the standard model and allows student to use their economic reasoning to consider a wider range of possible outcomes, make sense of recent empirical results, and be better prepared to follow public debates.

¹ I have used these tools in both principles and intermediate macroeconomics courses. In addition, I have used them in a macroeconomics elective aimed at seniors. In the latter case, I present the insights as background for an in-depth discussion of the relevant contemporary research.

² See, e.g., Autor (2015).

³ See Korinek and Stiglitz (2018) and Aghion, Jones, and Jones (2018) for recent discussions of the economics of artificial intelligence.

In addition to presenting the model of automation, I also discuss how the difference between the standard model and the new task-based model provides a useful starting point for a wider classroom discussion of labor markets, including skill-biased technical changes and immigration, two other important topics of public debate. The appendix includes several exercises that students can complete to gain a more comprehensive understanding of labor markets and think about the impacts of the rapidly changing technological environment in which we currently live. These include discussions of the economic impacts of large language models like ChatGPT and the consequences of a universal basic income (UBI), a policy that is commonly proposed as a means to mitigate the labor market consequences of new automation technologies.

In addition to improving course content, my experience suggests that the discussion of automation also increases students' interest in labor markets, especially students who are interested in working in the technology sector after graduation. To further demonstrate the importance of this topic, in my principles and intermediate macroeconomics courses, I include it as part of a larger module on inequality, a topic that many students are passionate about. The module mainly focuses on the fall in the labor share of income that started around 2000 (Jones, 2016). I cover the three leading theories about why this occurred: automation, monopoly power, and international trade (Grossman & Oberfield, 2022), all of which require extensions to the standard approach to labor markets used in undergraduate macroeconomics courses.

2. Technical Background

In this section, I present the mathematical background for the model. This discussion is aimed at instructors, rather than students, although it could be appropriate for more advanced undergraduate courses. Consider the production function

$$Y = A (B+L)^{1-\alpha} K^\alpha, \quad (1)$$

where Y is output, A is a measure of productivity, B is the quantity of “robots”⁴ or some other automation technology,⁵ L is human labor, K is standard physical capital, and $\alpha \in (0,1)$. This production function satisfies the usual properties of a neoclassical production function in that it has constant returns to scale in its three rival inputs (B , L , and K) and each of these inputs has a positive and diminishing marginal product.⁶

In production function (1), the defining feature of robots (B) is that they are perfect substitutes for labor (L). This is the essence of the recent macroeconomic literature: some technological advances lead to the creation of new types of capital goods (“robots” or “automation capital”) that are perfect substitutes for labor because they perform the same tasks as labor (e.g., Acemoglu & Restrepo, 2018a, 2019). This makes robots distinct from other types of capital -- like equipment, structures, or software -- that are more complementary to human labor. The discovery of this new type of capital good is a type of technical change, but it affects the economy differently than an increase in A , which is the usual way to capture technical change. I now turn to discussing the mathematical implications of the model, which makes this latter distinction clearer.

⁴ Here, I use “robots” following the recent literature, though in general the term does not necessarily imply a technology that is a perfect substitute for labor.

⁵ Here, robots are measured in terms of equivalent human workers. If robots were measured in a different unit, then it would be necessary to include another variable to account for differences in productivity between robots and humans. In this case, effective labor inputs would be $Z \cdot B + L$, where Z captures the relative productivity. Normalizing $Z=1$ is equivalent to choosing the units in which robots are measured.

⁶ The production function also satisfies the Inada conditions and essentiality in two inputs, K and $B+L$, where the latter is effective labor inputs.

The goal of this analysis is to present a simple supply and demand model of the labor market that can incorporate the impact of robots. From the firm's point of view, human labor and robots are indistinguishable: $B+L$ is the total quantity of "effective labor inputs". Assuming perfect competition, the demand curve for effective labor is

$$W/P = (1-\alpha) A(B+L)^{-\alpha} K^{\alpha}, \quad (2)$$

where W/P is the real wage. The real wage is equal to the marginal product of effective labor, the quantity shown on the right-hand side of the equation. If both human labor and robots are used in equilibrium, they must be paid the same rate. For simplicity, I use "wage" to denote the rental rate for robots.⁷ This expression highlights the key distinction between robots (B) and the standard notions of capital (K) or technology (A). An increase in B decreases the demand for human labor, while increases in K or A increase the demand for human labor.

To close the model, it is necessary to specify the supply of factors. As is standard in undergraduate models of labor markets, I assume that the supply of capital is fixed to focus on a partial equilibrium model of the labor market. Due to the labor-leisure trade-off, there is an upward sloping human labor supply curve,

$$W/P = S_L(L), \quad (3)$$

where S_L defines the supply curve and $S'_L > 0$. Since robots do not value leisure, I treat them like capital and assume that the supply is fixed, but this is not essential for anything that follows. Then, the joint supply curve for labor inputs is

$$W/P = S_L(L) + B, \quad (4)$$

where B is fixed. With this simple specification, (2) and (4) can be solved for W/P and L , given that B is fixed. In algebraic examples given to students, I usually focus on the case where S_L is linear. In this case, (2) and (4) can also be used to solve for W/P and the sum $(L+B)$. Then, plugging the equilibrium real wage back into (3) gives the equilibrium quantity of human workers (L). The model I present in the following section shows the solution to these equations graphically.

Before continuing, it is worth noting some caveats about the model presented above. First, the Cobb-Douglas functional form in (1) is chosen for convenience and is not essential for any of the analyses. In class, I often use linear demand curves to simplify the algebra for students. Second, I have only considered one type of general technology, A , which affects the marginal product of all inputs equally. The key insight is that A increases the demand for labor, while B decreases the demand for labor. The fact that an improvement in technology increases the demand for human labor holds for almost any type of factor-augmenting technology. One exception is that labor-augmenting technology can decrease the demand for labor when the elasticity of substitution between labor and capital is sufficiently low (Acemoglu & Restrepo, 2018b). The literature argues that this case is uninteresting because the low elasticity of substitution is inconsistent with data. Third, I have presented a static model. The academic literature is often focused on dynamics, which is beyond the scope of the pedagogical issues discussed here. Fourth, while this presentation is motivated by the intuition from the task-based literature, I do not present a full task-based model, because the mathematical complexity of such models is well beyond what is expected in undergraduate courses. My goal is to capture some of the key intuitions from this literature with minimal modifications to standard pedagogical tools.


⁷ Since the goal here is to introduce robots into a standard macroeconomic analysis, I assume that robots are rented by firms, which is the same way that capital is treated in standard models.

3. Automation in a Standard Supply and Demand Model

In this section, I present the model in a manner that is intended for students in courses that do not require calculus. As a starting point, take the standard supply and demand framework of the labor market (Figure 1). There is perfect competition. For notation, I use L as the number of employed workers, W as the nominal wage, and P as the price level. Workers face a tradeoff between labor and leisure. Willingness to work is increasing in wages. I use S_L to denote the supply curve for workers. Competition between firms implies that the height of the demand curve (D) is equal to the marginal product of labor, which I denote with MPL .


Asterisks (*) denote equilibrium values. The equilibrium occurs where the supply and demand curves intersect, which is where $MPL=W/P$. Intuitively, the equilibrium is where neither firms nor workers want to change behavior. Firms are maximizing profits, and workers are optimally choosing labor or leisure based on the equilibrium wage.

Figure 1: Equilibrium in the labor market



In standard undergraduate pedagogy, the impact of a technological advance is analyzed through an increase in the marginal product of labor. Graphically, this corresponds to an outward shift in labor demand, which increases wages and employment. This is shown in Figure 2. This analysis is consistent with predictions of standard economic theory. With a standard two-factor neoclassical production function, almost any increase in a factor-augmenting or factor-neutral technology leads to a higher marginal production of labor (Acemoglu & Restrepo, 2018b).


Figure 2: Standard Model of Technological Advance



Contrary to this prediction, recent empirical evidence suggests that some new technologies lead to lower wages and employment (e.g., Acemoglu & Restrepo, 2019, 2020). To capture this fact, an emerging academic literature often studies task-based models of automation (e.g., Zeira, 1998; Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2018a). The key insight from these models is that the production process within a firm often requires the completion of several different tasks. Some technological advances give machines a comparative advantage in a new set of tasks, displacing workers. The baseline task-based model captures this fact by considering the case where machines and workers are perfect substitutes within “automated” tasks. While the mathematical complexity of the full model is well beyond a standard undergraduate course, this basic insight about some new technologies being perfect substitutes can be easily captured in the standard supply and demand framework.

In Figure 3, I update the model to include a supply of “robots” that are perfect substitutes for human workers. Firms consider humans and robots to be interchangeable. So, there is still a single demand curve for human-like inputs to production. Since robots don’t value leisure, the supply of robot labor is perfectly inelastic.⁸ I denote the quantity of robots with $S_B = B$. Firms rent robots and pay only a rental rate.⁹ The joint supply curve of humans and robots is given by the curve $S = S_L + S_B$. This is just the S_L curve shifted to the right by B .

Figure 3: New Model of Technological Advance



The equilibrium occurs when the demand for human-like inputs is equal to their supply. This occurs where $S_L + S_B$ intersects $D = MPL$. The height of the intersection gives the equilibrium real wage, $(W/P)^*$, which is also the real rental rate for robots. The horizontal position of the intersection gives the joint quantity of workers and robots, $L^* + B$, hired by firms. Recalling that the labor supply curve gives the number of workers willing to work at a given wage, the horizontal position of S_L corresponding to $(W/P)^*$ gives the equilibrium supply of workers, L^* .

Comparing the equilibrium with robots to the equilibrium in Figure 1 (i.e., to the intersection of $D = MPL$ and S_L), we see that the introduction of robots has decreased both equilibrium wages and equilibrium employment. This is consistent with the empirical evidence

⁸ In reality, of course, robots need maintenance and may have an upward-sloping supply curve. By abstracting from maintenance, I am treating “robots” in the same way that standard macroeconomic models and pedagogy treat traditional capital (K).

⁹ This follows the standard treatment of traditional capital in macroeconomic models.

on automation and wages cited above. In addition, total output has increased, because there are more inputs to production. Together, the increase in output and decrease in wages and employment imply that the labor share of income has fallen, providing a possible explanation for the observed macroeconomic trends.

The decrease in wages should be very intuitive for students. There has been an outward shift in the supply of human-like inputs. Following standard supply and demand logic, this leads to a lower price and a higher quantity. The simple framework presented here differs from the standard supply and demand model in that we happen to care not only about the total quantity of human-like inputs hired by firms but also about the disaggregation between humans and robots. To determine what happens to humans, we simply need to check the human labor supply curve.

The above discussion shows how a simple adjustment to the supply and demand model of labor markets can help students understand the consequences of automation technologies. This is accomplished by considering two supply curves: a joint supply curve for humans and workers, as well as the standard supply curve for humans only. While this makes the model slightly more complicated, it is about as complicated as other models taught in classes without calculus. For example, students often encounter multiple supply curves when thinking about international capital flows, externalities, or taxation. With this small change, students will better understand the perspectives of academic economists and be better positioned to follow public debates about the labor market impacts of new technologies.

4. Further Discussion

The previous section focused on how to describe the equilibrium in the augmented supply and demand model and on how to compare the equilibrium to the more standard outcomes without robots. The latter is essentially a comparative statics exercise of increasing the supply of robots from 0 to B . As discussed in section 5, it is straightforward to use the model to conduct any other standard comparative statics exercise used to study labor markets in undergraduate courses.

In this section, I consider how this simple model and underlying logic can be used as a springboard to discuss other important topics in labor markets, such as skill-biased technical change and immigration. In doing so, I demonstrate how the task-based perspective on labor markets allows students to utilize a very simple framework to analyze important public policy issues in more depth than is possible with standard tools.

4.1 Tasks and Worker Heterogeneity

The augmented supply and demand model presented in the previous section highlights the intuition of *tasks*, a theoretical concept that plays a significant role in recent advances in the academic literature. The intuition of tasks is also a natural setting in which to think about heterogeneity between workers. Understanding this heterogeneity, in turn, is a prerequisite for students to be able to carefully analyze skill-biased technical change and immigration in ways that are consistent with cutting-edge academic research and empirical evidence.

To illustrate the importance of heterogeneous workers to students, I often focus on a simple example of a technology company with two types of tasks: there are software engineers who write computer code and custodians who keep the office clean and safe. Both tasks contribute to the output (software) of the firm. The engineers directly produce the software, and the custodians contribute indirectly by increasing the productivity of the software engineers, who can work more efficiently in a clean, safe, and well-maintained office environment.

Now, we can ask: “How does the marginal productivity of the custodians change

when the number of software engineers increases?" There is good reason to believe that their productivity has increased because there are more software engineers whose productivity is affected by custodial tasks like taking out the trash and vacuuming the floor. The same argument works in reverse: an increase in the number of custodians leads to a better work environment and increases the productivity of each software engineer. Conversely, due to the usual story about decreasing returns, we expect that the marginal productivity of each software engineer decreases when a new engineer is added.

4.2 Skill Biased Technical Change and Inequality

Once students start thinking about workers who perform different tasks, it is straightforward to think about skill-biased technical change. A new technology can perform a given set of tasks. Workers who perform these tasks will find themselves in the situation highlighted in Figure 3. Their wages and employment levels will decrease. On the other hand, workers who perform complementary tasks will see their marginal product increase, leading to the outcome depicted in Figure 2, which indicates higher wages and employment.

So, which workers are vulnerable to the introduction of new automation technologies? This is a difficult question, and the answer seems to evolve over time. A piece by David Autor in the *Journal of Economic Perspectives (JEP)* provides a clear and informative overview that is accessible to undergraduate students (Autor, 2015). A key insight of the piece is the role of *routine* tasks. Automation technologies are particularly good at performing tasks that can be codified in a set of simple, repetitive rules. These tasks tend to be performed by workers in the middle of the income distribution. On the other hand, automation technologies are not good at tasks that involve tacit knowledge and abstract thinking. As a result, they have trouble replacing low-income workers who do manual tasks like cleaning and food preparation. They also have trouble replacing high-income workers who are engaged in abstract planning and managerial tasks.

The piece by Autor was written in 2015, before the recent advances in large language models (LLMs), like Chat GPT.¹⁰ Recent advances in LLMs and other forms of artificial intelligence (AI) may well change the type of tasks that are vulnerable to automation. As discussed in the next section, the changing nature of task vulnerability is an interesting topic for class discussions or writing assignments.

The comparison of Figures 2 and 3 is also an intuitive lead-in to classroom discussions about how to address the inequality engendered by automation technologies and the dynamics of that inequality over time. The tasks that a given worker can complete are not fixed. With education or job training, workers might be able to move from Figure 3 to Figure 2. In addition, workers who chose to invest in skills before the introduction of the new technology are more likely to end up in Figure 3, compared to workers who can choose skill after the introduction, suggesting the short-run and long-run effects of automation technologies can be quite different. This distinction is stressed in the academic literature (e.g., Acemoglu and Restrepo, 2018).

4.3 The Impact of Immigration on Domestic Workers

The intuition of tasks is also a natural setting to think about the economic impacts of immigration on workers' wages already residing in a country ("domestic workers"). I have found that the model presented in section 3 can be combined with a piece by Giovanni Peri in the *JEP* to give students a succinct overview of economists' views of immigration (Peri, 2016). I provide a summary of that discussion below.

¹⁰ For a discussion of large language models aimed at economists, see <https://bcf.princeton.edu/events/kevin-bryan-on-chatgpt-a-research-tool-in-economics/>.

The standard labor market model used in undergraduate education focuses on a homogenous pool of labor. In this setting, immigration pushes out the supply curve, increasing total employment and lowering total wages. This is consistent with how academic economists studied labor markets in the 1980s and 1990s, but inconsistent with more recent theory and empirical evidence.

As in the case of skill-biased technical change, immigration has different impacts on different workers already residing in a country. The experience of domestic workers who perform the same tasks as immigrant workers is roughly captured in Figure 3. Here, L would be domestic workers performing similar skills, and B would be immigrant workers.¹¹ The experience of domestic workers who perform different tasks than immigrant workers is then roughly captured in Figure 2. In this case, the outcome for immigrant workers is not shown.

So, which domestic workers perform similar tasks to immigrant workers? There is disagreement among economists on this question and it makes an interesting topic for writing assignments or in-class discussions. As suggested by the previous section, education is one important determinant of the task content of work. Domestic workers without a high school degree are mostly likely to perform the same tasks as immigrant workers without a high school degree, while domestic workers with advanced engineering degrees are mostly likely to perform the same tasks as immigrant workers with advanced engineering degrees, and so on. Even within education categories, however, the task content of work can differ considerably. Domestic workers tend to specialize in tasks involving communication, perhaps because of the comparative advantage of language skills. For example, domestic workers without a college degree may be more likely to work in customer-facing retail positions, while immigrant workers are more likely to be in positions where manual labor is important. Domestic workers with a college degree are more likely to be in managerial positions, while immigrant workers are more likely to work in positions that place a high value on mathematical and analytical skills. In either case, an increase in immigrant workers may boost the productivity of domestic workers with similar skill levels, and an increase in immigrant workers may be most likely to decrease the wages of previous immigrant workers with similar skill levels.

5. Exercises and Activities

In this section, I discuss several exercises that students can complete to deepen their understanding of the model and apply the intuition to a broad range of questions.

Finding an equilibrium: The discussion in section 3 focuses on a graphical presentation of the supply and demand model with automation. Of course, it is also possible to study the model algebraically, which can help reinforce the key concepts. Appendix B presents a sample homework problem that walks students through algebra.

Comparative statics: As with any supply and demand model, comparative statics in the automation model can be used to study a wide range of economic policies and exogenous shocks. In class, I like to investigate the consequences of a universal basic income (UBI), which is a commonly proposed policy to combat the effects of automation. This translates to an inward shift in s_L , increasing the wage and further decreasing employment. One downside of the model is that conducting comparative statics analyses following shifts in supply requires moving multiple curves. This is similar to shifting multiple curves in a supply and demand model with externalities and/or taxes (e.g., after a change in private marginal costs).

Extensions: As noted in section 4, a key benefit of the supply and demand model of automation is that it is a useful starting point for discussing a wide range of issues in

¹¹ Figure 3 would then technically imply that the labor supply of immigrant workers is perfectly inelastic, but it could be updated to make $s_L + s_B$ flatter than s_L , reflecting the labor-leisure trade-off for immigrant workers.

the macroeconomics of labor markets. As a result, there are open-ended questions that can serve as effective prompts for a class or small group discussions, or as short-answer questions for problem sets and exams. The appendix includes some examples focused on the potential economic impacts of LLMs like ChatGPT.

Small group discussions: Given that the model connects to ongoing policy debates, it is also a natural starting point for small group discussions. In my classes, I have small groups of students read a recent academic article that can be understood using the model. They discuss the paper with me in a small group setting and then present the lessons they learned from their small group discussion to the rest of the class, usually in a recorded video. Groups are chosen based on a survey of student interests distributed at the beginning of the semester. This activity has at least three benefits for students. First, it allows them to dive into the topic in greater depth than can be covered in class. Second, it shows them that the models they learn in class are used by economists to think about current policy debates. Third, it reinforces key communication and presentation skills. I find this activity tends to increase student engagement with and interest in the material. I use articles from the *Journal of Economic Perspectives (JEP)*, which tend to be accessible to principles and intermediate macroeconomics students, even if they do not necessarily have the background to understand all of the equations or statistical arguments. I have used the model presented in this paper to discuss three *JEP* articles: Autor (2015), Acemoglu and Restrepo (2019), and Fort, Pierce, and Schott (2018). The first two articles are explicitly framed using the task-based framework. The first discusses the relationship between automation, skill-biased technical change, and the hollowing out of the labor income distribution. The second covers the effect of automation on employment and the labor share of income. The third article compares different causes for the decline in US manufacturing employment, focusing on the relative impacts of trade and automation.

6. Conclusion

One of the main goals of undergraduate economics education is to help students understand the world around them. In recent years, there have been significant technological advances that may have significant impacts on labor markets and inequality. Industrial robots and large language models are two obvious examples. Unfortunately, the standard approach to discussing labor markets in macroeconomics courses doesn't give students all of the relevant tools they might need to analyze these trends.

In this paper, I show how a small modification to the standard supply and demand framework makes it possible for students to analyze the introduction of automation technologies. I also discuss how the model is related to a task-based view of labor markets. This approach is both related to the academic literature and helpful for discussing skill-biased technical change and immigration. The simplicity of the modified model makes it possible for students without a background in calculus to investigate these topics, which also play a large role in public debates about economic policy, especially those surrounding trends in inequality and the labor share of income.

This paper is motivated in part by a desire to close the gap between academic macroeconomics and the macroeconomics taught in undergraduate classes. A downside of teaching students about cutting-edge economic ideas is that these ideas are not settled. The purpose of this article is not to argue that the new approach to modeling labor markets will necessarily turn out to be a more accurate description of the real world. After all, there have been many occasions in the past where concerns about the employment effects of new technologies did not materialize (Autor, 2015; Mokyr, Vickers, & Ziebarth, 2015). Instead, the goal of this paper is to ensure that students are exposed to a wide range of ideas that represent different perspectives reflected in cutting-edge economic research and related policy discussions.

References

- Acemoglu, D. 2013. Economic growth and development in the undergraduate curriculum. *Journal of Economic Education*, 44(2), 169-177. DOI: [10.1080/00220485.2013.770344](https://doi.org/10.1080/00220485.2013.770344)
- Acemoglu, D., & Autor, D. 2011. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, vol. 4, 1043-1171. Elsevier. DOI: [10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., & Restrepo, P. 2018a. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542. DOI: [10.1257/aer.20160696](https://doi.org/10.1257/aer.20160696)
- Acemoglu, D., & Restrepo, P. 2018b. Modeling automation. *AEA Papers and Proceedings* 108, 48-53. DOI: [10.1257/pandp.20181020](https://doi.org/10.1257/pandp.20181020)
- Acemoglu, D., & Restrepo, P. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30. DOI: [10.1257/jep.33.2.3](https://doi.org/10.1257/jep.33.2.3)
- Acemoglu, D., & Restrepo, P. 2020. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188-2244. DOI: [10.1086/705716](https://doi.org/10.1086/705716)
- Aghion, P., Jones, B. F., & Jones, C. I. 2018. Artificial intelligence and economic growth. In *The Economics of Artificial Intelligence: An Agenda* (237-282). University of Chicago Press.
- Autor, D. H. 2015. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30. DOI: [10.1257/jep.29.3.3](https://doi.org/10.1257/jep.29.3.3)
- Fort, T. C., Pierce J.R., & Schott, P. K. 2018. New perspectives on the decline of US manufacturing employment. *Journal of Economic Perspectives*, 32(2), 47-72. DOI: [10.1257/jep.32.2.47](https://doi.org/10.1257/jep.32.2.47)
- Grossman, G. M., & Oberfield, E. 2022. The elusive explanation for the declining labor share. *Annual Review of Economics*, 14, 93-124. DOI: [10.1146/annurev-economics-080921-103046](https://doi.org/10.1146/annurev-economics-080921-103046)
- Hoyt, G. M., & McGoldrick, K. M. 2019. 50 years of economic instruction in the Journal of Economic Education. *Journal of Economic Education*, 50(2), 168-195. DOI: [10.1080/00220485.2019.1582388](https://doi.org/10.1080/00220485.2019.1582388)
- Ihrig, J. & Wolla, S. 2022 Let's close the gap: Updating the textbook treatment of monetary policy. *Journal of Economic Education*, 53(2), 232-249. DOI: [10.1080/00220485.2022.2075509](https://doi.org/10.1080/00220485.2022.2075509)
- Ihrig, J., & Wolla, S. 2023. Teaching monetary policy with ample reserves. *Journal of Economics Teaching*, 8(2), 114-127. DOI: [10.58311/jeconteach/6f8dff11179466b2d32da497c5465b9cff2774ac](https://doi.org/10.58311/jeconteach/6f8dff11179466b2d32da497c5465b9cff2774ac)
- Jones, C. I. 2016. The facts of economic growth. In *Handbook of macroeconomics*, vol. 2, 3-69. Elsevier. DOI: [10.1016/bs.hesmac.2016.03.002](https://doi.org/10.1016/bs.hesmac.2016.03.002).
- Jones, C. I., & Romer, P. M. 2010. The new Kaldor facts: ideas, institutions, population, and human capital. *American Economic Journal: Macroeconomics*, 2(1), 224-245. DOI: [10.1257/mac.2.1.224](https://doi.org/10.1257/mac.2.1.224)
- Korinek, A., & Stiglitz, J. E. 2018. Artificial intelligence and its implications for income distribution and unemployment. In *The Economics of Artificial Intelligence: An Agenda*, 349-390. University of Chicago Press. DOI: [10.7208/9780226613475-016](https://doi.org/10.7208/9780226613475-016)

Mokyr, J., Vickers, C., & Ziebarth, N. L. 2015. The history of technological anxiety and the future of economic growth: Is this time different? *Journal of Economic Perspectives*, 29(3), 31-50. DOI: [10.1257/jep.29.3.31](https://doi.org/10.1257/jep.29.3.31)

Peri, G. 2016. Immigrants, productivity, and labor markets. *Journal of Economic Perspectives*, 30(4), 3-30. DOI: [10.1257/jep.30.4.3](https://doi.org/10.1257/jep.30.4.3)

Solis-Garcia, M. 2018. The macro pedagogy debate: Teaching DSGE to undergraduates symposium: Yes we can! Teaching DSGE models to undergraduate students. *Journal of Economic Education*, 49(3), 226-236. DOI: [10.1080/00220485.2018.1464987](https://doi.org/10.1080/00220485.2018.1464987)

Solow, R.M. 1956. A contribution to the theory of economic growth. *Quarterly Journal of Economics*, 70(1), 65-94. DOI: [10.2307/1884513](https://doi.org/10.2307/1884513)

Zeira, J. 1998. Workers, machines, and economic growth. *Quarterly Journal of Economics*, 113(4), 1091-1117. DOI: [10.1162/003355398555847](https://doi.org/10.1162/003355398555847)

Appendix B: Example Questions and Solutions

In this appendix, I present some exercises that use the model and ideas discussed in this paper. Section 4 discusses the pedagogical goals behind the questions. Sample solutions are in blue and offset from the questions. Notes in brackets [] are for the reader of this paper and would not be given to students.

Question 1: Finding an equilibrium with algebra

Consider a labor market without unemployment. Suppose that the number of people willing to work in a given economy is given by $L = 20 + (W/P)$, where L is people, W is the nominal wage, and P is the price level. There are also 10 robots in the economy (i.e., $B=10$). As in class, robots and people are interchangeable in production. The marginal product of labor is given by $MPL = 60 - (L + B)$.

The diagram makes it easier to keep track of everything.

[Reproduce figure 1c here]

- a. What is the equilibrium wage in this economy?

Here are two (equivalent) ways of solving the problem.

(1) We know that $W/P = MPL$ on the demand curve. The demand curve is $W/P = 60 - (L + B)$. For the supply curve, we add the supply of workers and the supply of robots: $E+B = (20 + W/P) + 10 = 30 + W/P$. Rearranging, $W/P = (L+B) - 30$. Setting supply = demand:

$$60 - (L + B) = (L + B) - 30 \quad L + B = 45.$$

From the demand curve, **$W/P = 60 - (L + B) = 15$.**

(2) We know demand is $W/P = 60 - (L+B)$ and $B = 10$, so $W/P = 50 - L$. We know the supply of workers is given as $L = 20 + W/P$. So, $W/P = L - 20$.

Now, set the two expressions equal to each other: $50 - L = L - 20 \quad 2L = 70 \quad L = 35$. Then, plug in L for the MPL equation so $60 - 35 - 10 = 15$.

- b. What is the equilibrium marginal product of labor?

$MPL = W/P = 15$.

- c. What is the equilibrium level of employment?


Note that $B = 10$. So, $L + 10 = 45 \quad L = 35$. You could also find this same result by plugging $W/P = 15$ into the supply curve for human workers.

Question 2: Universal Basic Income as a Policy Response to Automation

Using the supply and demand model of automation, show how a UBI affects each of the following.

- a. The equilibrium wage of human workers

For all of the questions, it will help to draw the picture, which is below.



The supply curve for workers is upward-sloping because of the labor-leisure trade-off. A UBI pays individuals whether or not they work. This will induce an income effect, but not a substitution effect. In other words, the demand for leisure increases (workers want more leisure at a given wage). This shifts both supply curves to the left by equal amounts. This shift is equal, because robots don't have a labor-leisure trade-off, implying that S_B is unaffected.

The new equilibrium wage is determined by the intersection of the new joint supply curve ($S_{L,2} + S_B$) and the original demand curve. **The equilibrium wage increases** relative to the previous equilibrium. This is consistent with an inward supply shock in any supply and demand model.

b. Equilibrium employment

The new equilibrium level of employment is given by the intersection of the human supply curve ($S_{L,2}$) and the demand curve. **Equilibrium employment decreases**, which is unsurprising given that we have had an inward supply shock to human workers.

c. Total output

There hasn't been any change in the production technology in this economy, but the quantity of inputs (specifically, workers) has decreased. So, **output has fallen**.

Question 3: Further Implications

Recently, there have been tremendous technological advances in LLMs like ChatGPT. As with any new technology, this has sparked a lot of public discussion about the economic impacts of this technological advance.

Answer each of the following questions about the economic consequences of ChatGPT in three or four sentences. I am looking for some relevant intuition and there are no right or wrong answers. No equations are necessary. You may refer to the figures from the class notes [Note: these correspond to Figures 1-3 in this paper.]

a. Briefly explain how ChatGPT might affect the labor share of income in the future.

Sample Answer: LLMs like ChatGPT can perform several tasks previously only possible with human labor. For example, they can write simple computer code or summarize a long document almost instantaneously. In this sense, they are much like "robots" from our automation model, and they increase the supply of human-like production inputs. As shown in Figure 3, this

increases the total quantity of output, but decreases both wages and employment, leading to a lower labor share of income.

- b. Briefly explain how ChatGPT might affect inequality between different groups of workers in the future.

Sample Answer: While LLMs like ChatGPT can perform some tasks in a manner very similar to humans (e.g., summarizing a long document), they cannot perform other tasks, like those involving manual labor. Since different types of tasks tend to be complements, manual task workers will experience an outward shift in demand, as shown in Figure 2. The example of manual labor suggests that ChatGPT might compress the distribution of labor income (i.e., increase wages at the bottom of the distribution and decrease them at the top), but ChatGPT also cannot perform managerial tasks, implying that it might have a hollowing out effect much like other automation technologies we have discussed. So, using a law firm as an example, ChatGPT might perform tasks similar to a paralegal, but not similar to building maintenance or a partner who is overseeing the strategy behind a case.

- c. Suppose a policymaker asked you to design a policy that would limit the impacts of ChatGPT on inequality. What would you recommend?

Sample Answer: The skills of an individual are not fixed, so I would recommend subsidizing training programs for workers affected by LLMs that help them develop skills complementary to those of the LLMs. The programs would focus on skills distinct from LLMs (like communication and management skills) or on skills that help people better utilize LLMs. As we discussed in class, the number of job openings for “prompt engineers” (people who design effective instructions for LLMs) has increased rapidly. It seems plausible that a worker replaced by an LLM might have a comparative advantage in thinking about how LLMs can be used more effectively in that job.