NBER WORKING PAPER SERIES

AI ADOPTION IN A COMPETITIVE MARKET

Joshua S. Gans

Working Paper 29996
http://www.nber.org/papers/w29996

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

Thanks to Daron Acemoglu, Avi Goldfarb and Richard Holden for helpful discussions. All errors remain my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

AI Adoption in a Competitive Market
Joshua S. Gans
NBER Working Paper No. 29996
April 2022
JEL No. D21,D41,D81,O31

ABSTRACT

Economists have often viewed the adoption of artificial intelligence (AI) as a standard process innovation where we expect that efficiency will drive adoption in competitive markets. This paper models AI based on recent advances in machine learning that allow firms to engage in better prediction. Using prediction of demand, it is demonstrated that AI adoption is a complement to variable inputs whose levels are directly altered by predictions and use is economised by them (that is, labour). It is shown that, in a competitive market, this increases the short-run elasticity of supply and may or may not increase average equilibrium prices. There are generically externalities in adoption with this reducing the profits of non-adoptees when variable inputs are important and increasing them otherwise. Thus, AI does not operate as a standard process innovation and its adoption may confer positive externalities on non-adopting firms. In the long-run, AI adoption is shown to generally lower prices and raise consumer surplus in competitive markets.

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

As a general rule, technologies that improve the efficiency of production tend to interact with competitive effects to encourage their adoption. For instance, a process innovation that lowers marginal costs or raises factor productivity, will be adopted by competitive firms, in part, because the profits of non-adoptees fall as others adopt; that is, adoptees create negative external effects on non-adoptees. As more firms adopt, prices fall as do the profits of those who do not adopt.

This paper demonstrates that this general rule does not generally apply to artificial intelligence (AI). Recent research has viewed AI as equivalent to task automation and hence, a standard process innovation.[1] The tendency, however, to equate AI adoption with automation does, however, mask the fact that recent advances in AI are all advances in machine learning that permits lower cost and/or better quality predictions.[2] Such advances in prediction can enable more opportunities for automation by machines – by combining that prediction with machine actions such as might arise for self-driving vehicles – but, at the same time, this has caused some to suggest that AI advances could be directed or redirected to augment rather than replace human labour (Markoff (2016), Acemoglu & Restrepo (2020b), Brynjolfsson (2022)).

In this paper, I take the notion that AI is prediction seriously in order to properly characterise the type of technological change it represents and what would guide its adoption in a competitive market. What predictions do is allow firms to match decisions, such as output and employment, with the predicted state. Thus, while a firm without such predictions may make decisions based on the average external state, AI prediction allows those decisions to vary. Thus, while output might be relatively stable when there is no prediction, the availability of a prediction may cause firms to increase or reduce output accordingly. While this represents an improvement in efficiency, it will be shown here that because some of the efficiency gain comes from reducing output, it may be that the external effect of AI adoption on other firms is positive rather than negative. Rather than adoption spurring others to adopt,

[1]Recent work on the impact of artificial intelligence (AI) has classified such as advances as primarily involving increased automation. For instance, Acemoglu & Restrepo (2018) view production functions as the aggregate outcome from the performance of a number of tasks. They distinguish between technological change that augments labour productivity in a task directly from that which allows capital to replace labour in a task or augments capital productivity in a task. They show that "the implications of automation and AI are generally very different from those of labor-augmenting advances" (Acemoglu & Restrepo (2018), p.212). This framework has guided much of the empirical work in trying to analyse and forecast the potential effect of AI on jobs and productivity. (See, for example, Brynjolfsson & Mitchell (2017), Felten et al. (2018), Susskind (2020), Das et al. (2020), Felten et al. (2021), Acemoglu & Restrepo (2020a), Acemoglu et al. (2020), Acemoglu & Restrepo (2021)).

[2]See Agrawal et al. (2018), Agrawal et al. (2019a), Agrawal et al. (2019b).

the opposite may be the case with the returns to adoption falling.

To model this, I focus here on the canonical environment where competitive firms face uncertainty regarding market demand for their products. My starting point is an observation by Nelson (1961) that, when facing uncertainty about demand, a firm's behaviour critically depends on how well it can make predictions regarding that demand. The value of such predictions is that the firm can make short-run adjustments to output that, when marginal costs are rising, allow it to more efficiently match output choices to realised market prices. This has important implications for how we characterise AI as a technological advance. In particular, at its heart, AI prediction is useful only insofar as it can change decisions. In the context of a price-taking firm, the relevant decisions involve the employment of factors of production to generate output. If the predictions are useful in the short-run, then we can refine the relevant decisions even further to focus on factor employment decisions themselves made in the short-run. As will be shown, this implies that AI fundamentally involves augmenting the productivity of factors that are variable in the short-run. That is, in this way *AI augments the productivity of variable rather than fixed inputs.* In the case of many firms this means that they augment labour directly rather than capital which is typically fixed in the short-run.

Nelson (1961) explored the value of adopting predictions by such firms but did not embed it in an equilibrium model of competitive market outcomes. He, therefore, did not examine the interaction between the profits of adopters and non-adopters. Here I provide that analysis while abstracting away from some the imperfections in predictive accuracy that were the focus of the older literature (see, in particular, Pashigian (1974)). Several results emerge. First, the above intuition that AI is labour augmenting is confirmed. Second, the value of AI to a firm comes from the variance in market prices caused by variation in demand. Third, as more firms adopt AI, this changes the variance in market prices – reducing it – as those firms adjust short-term output in a way the mitigates price variability. Fourth, AI adoption can have positive or negative external effects on the profits of firms who do not adopt AI. This may spur or diminish the diffusion of AI in a competitive market. Fifth, AI adoption can, in some cases, have a negative impact on consumer surplus even in a competitive market. Sixth, when firms interact in factor markets (such as labour) prior to making output decisions, prices in those markets can signal prediction outcomes. This informational spillover further reduces the incentives to adopt AI. Finally, in the long-run, when capital can adjust, market prices are, on average lower, and quantity is, on average, higher with AI adoption in the industry than without it.

In what follows, I begin with a short-run analysis before turning to interactions between product and factor markets and then to a long-run treatment. The main conclusion reached is that what AI does matters for how it enters into the production function of firms and that in competitive markets, how this translates into pricing and welfare outcomes can be subtle. Importantly, there are reasons to suppose that competitive forces will not necessarily spur AI adoption.

2 Short-Run Analysis

There is a continuum of firms, $i \in [0, 1]$ in a market. Firms have an identical production function $f(L_i, K_i)$ that is non-decreasing and convex in the quantity of labour, L_i and capital, K_i. The wage rate is w with labour supply perfectly elastic to the industry. The labour choice is variable in the short-run. In the long-run, capital can be altered at a rental rate, r, per unit of capital chosen. Initially, as factor prices are not the focus, w is set equal to 1 for expositional simplicity.

Firms produce perfectly substitutable, non-storable products to consumers whose market demand is given by:

$$Q = D(P, \theta) \tag{1}$$

where $D(.)$ is non-increasing in P and non-decreasing in $\theta \in \mathbb{R}_+$ where $\theta \sim \Phi(\theta)$. Upon the realisation of θ, the equilibrium market price, $\hat{P}$, is determined by market clearing; i.e.,

$$D(\hat{P}, \theta) = \int_0^1 f(L_i, K_i)di$$

Firms must choose and pay labour prior to the resolution of uncertainty. In this section, capital is assumed to be fixed at a common level for each firm.

In what follows, we will rely on the following functional forms.

$$D(P, \theta) = \theta - P$$

$$f(L_i) = L_i^{\alpha} K_i^{1-\alpha}$$

for $\alpha < 1$. To make exposition here simple, it assumed that for each i, capital employed is fixed at $K = 1$.[3] Assuming that all firms employ the same amount of labour, L, this implies that the market price is:

$$\hat{P}(\theta) = \theta - L^{\alpha}$$

As Nelson (1961) shows, firms will choose their optimal quantity taking expected price

[3] In a subsequent section, firms choose their long-run level of capital.

$E[\hat{P}]$ as given. Each firm solves:

$$\max_{L_i} E[\hat{P}]L_i^\alpha - L_i \implies \hat{L} = \left(\alpha E[\hat{P}]\right)^{\frac{1}{1-\alpha}}$$

Thus, in equilibrium, output will be:

$$\hat{Q} = (\alpha E[P])^{\frac{\alpha}{1-\alpha}} = \alpha^{\frac{\alpha}{1-\alpha}} \left(\int \theta d\Phi(\theta) - \hat{Q}\right)^{\frac{\alpha}{1-\alpha}}$$

with price, $\hat{P} = \theta - \hat{Q}$.

This means that expected profits will be:

$$E[\hat{\pi}] = (1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} E[\hat{P}]^{\frac{1}{1-\alpha}}$$

2.1 AI Adoption by One Firm

If θ was known with certainty by one firm, then $\hat{P}(\theta) = \theta - (\alpha E[P])^{\frac{\alpha}{1-\alpha}}$, as output is determined by the firms who do not know θ. So:

$$\pi(\theta) = (1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} \hat{P}(\theta)^{\frac{1}{1-\alpha}}$$

This implies that:

$$E[\pi(\theta)] = (1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} \int \hat{P}(\theta)^{\frac{1}{1-\alpha}} d\Phi(\theta)$$

Thus, starting from a position where no firm can predict demand, the willingness to pay of a firm for a perfect prediction is:

$$E[\pi(\theta)] - E[\hat{\pi}] = (1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} \left(E[\hat{P}^{\frac{1}{1-\alpha}}] - E[\hat{P}]^{\frac{1}{1-\alpha}}\right) > 0$$

where the sign comes from Jensen's Inequality and the fact that one infinitesimal firm having information will not change the expected price distribution in equilibrium.

It is critical to note here that the adoption of AI is enhancing the productivity the firm is able to generate from employing labour. In this respect, it is labour-augmenting technological change. However, because it augments labour as a factor that can be adjusted to maximise productivity, this does not necessarily translate into a high quantity of labour employed, on average.

To see this, note that the expected output of that firm is:

$$E[\hat{q}(\theta)] = \alpha^{\frac{\alpha}{1-\alpha}} \int \hat{P}(\theta)^{\frac{\alpha}{1-\alpha}} d\Phi(\theta)$$

It is useful to observe the following (again using Jensen's inequality):

$$\alpha \geq (<) \tfrac{1}{2} \implies E[\hat{q}(\theta)] \geq (<) \hat{Q}$$

That is, the faster the marginal product of labour falls with output (i.e., the lower is α), the average output chosen by a firm with AI will be lower, including, if $\alpha < \frac{1}{2}$, less than $\hat{Q}$. Thus, while AI enhanced profits, this may be associated with higher or lower output on average.

Example 1 *Suppose that $\theta = 1$ with probability ρ and $\theta = 2$ with probability $1-\rho$. Also set $\alpha = \frac{1}{2}$. Then note that $\hat{L}_i = (\frac{1}{2}E[\hat{P}])^2$ while $\hat{Q} = \frac{1}{2}E[\hat{P}]$. Then $E[\hat{P}] = (\rho + (1-\rho)2 - \frac{1}{2}E[\hat{P}] \implies E[\hat{P}] = \frac{2}{3}(2-\rho)$. Thus, $\hat{L} = (\frac{2-\rho}{3})^2$ and $\hat{Q} = \frac{2-\rho}{3}$. If $\theta = 1$, $\hat{P}(1) = 1 - \frac{2-\rho}{3}$ and if $\theta = 2$, $\hat{P}(2) = 2 - \frac{2-\rho}{3}$. This implies that expected profits are: $E[\hat{\pi}] = E[\hat{P}]\hat{Q} - \hat{L} = 2\left(\frac{2-\rho}{3}\right)^2 - \left(\frac{2-\rho}{3}\right)^2 = \frac{1}{9}(2-\rho)^2$. These pricing and quantity outcomes are depicted in Figure 1.*

If one firm adopts AI, this will leave market price unchanged but that firm can forecast $\hat{P}$ and adjust labour accordingly. If $\theta = 1$, $\hat{L}(1) = (\frac{1}{2}(1 - \frac{2-\rho}{3}))^2$ and if $\theta = 2$, $\hat{L}(2) = (\frac{1}{2}(2 - \frac{2-\rho}{3}))^2$ resulting in $\hat{q}(1)$. That is, the firm will adjust quantity based on the market price. Thus, expected profits become:

$$E[\pi(\theta)] = \rho\left(E[\hat{P}(1)]\sqrt{\hat{L}(1)} - \hat{L}(1)\right) + (1-\rho)\left(E[\hat{P}(2)]\sqrt{\hat{L}(2)} - \hat{L}(2)\right)$$

$$= \tfrac{1}{4}\left(\rho(1 - \tfrac{2-\rho}{3}))^2 + (1-\rho)(2 - \tfrac{2-\rho}{3}))^2\right) = \tfrac{4}{9} - \tfrac{1}{36}\rho(5\rho + 7)$$

As expected as $\alpha = \frac{1}{2}$, the expected output of a firm with AI is the same as one with AI but because that output is matched to the state, expected profits are higher.[4] It is easy to see that the return to AI adoption, $E[\pi(\theta)] - E[\hat{\pi}] = \frac{1}{4}\rho(1-\rho)$ is maximised when $\rho = \frac{1}{2}$; that is, when there is the greatest degree of uncertainty over the value of θ.

2.2 AI Adoption by Many Firms

Now suppose that a fraction, k of firms adopt AI. Let $q(\theta)$ be the symmetric output of those firms contingent on θ and q be the output – the same across states – of the $1-k$ who do

[4] A result demonstrated in Nelson (1961).

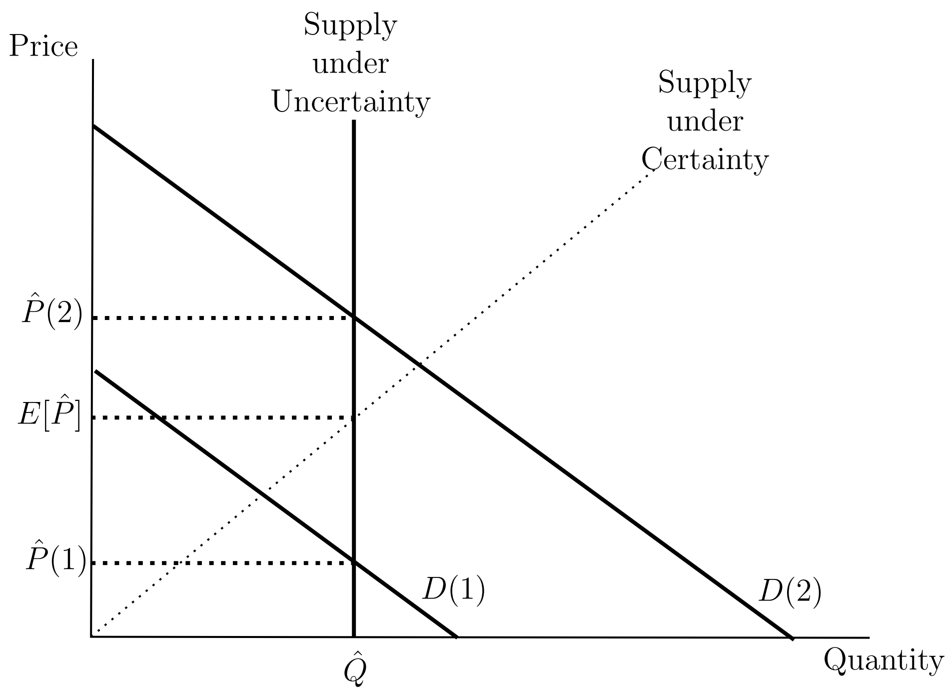

Figure 1: Market Equilibrium without AI Adoption

not adopt AI. Then

$$\hat{P}(\theta) = \theta - (1-k)q - kq(\theta)$$

Note that:

$$\hat{q} = (\alpha E[\hat{P}])^{\frac{\alpha}{1-\alpha}} = \alpha^{\frac{\alpha}{1-\alpha}} \left(\int \theta d\Phi(\theta) - k \int \hat{q}(\theta) d\Phi(\theta) - (1-k)\hat{q} \right)^{\frac{\alpha}{1-\alpha}}$$

$$\hat{q}(\theta) = (\alpha \hat{P}(\theta))^{\frac{\alpha}{1-\alpha}} = \alpha^{\frac{\alpha}{1-\alpha}} (\theta - k\hat{q}(\theta) - (1-k)\hat{q})^{\frac{\alpha}{1-\alpha}}$$

The first result of interest is to compare expected prices and quantities following the adoption of AI by k firms with those outcomes when no firm adopts AI.

Proposition 1 *If* $\alpha \geq (<)\frac{1}{2}$, $E[\hat{Q}(k)] \geq (<) E[\hat{Q}(0)]$ *and* $E[\hat{P}(k)] \leq (>) E[\hat{P}(0)]$ *for* $k > 0$.

Proof. Note that $\hat{q}(k)$ is increasing in $E[\hat{P}]$. Start from a point where $\hat{Q}(0) = \hat{q}(k)$. This means that $E[\hat{Q}(k)] = (1-k)\hat{q} + k \int \hat{q}(\theta) d\Phi(\theta) = \hat{Q}(0) + k \int (\hat{q}(\theta) - \hat{Q}(0)) d\Phi(\theta)$. If it assumed that the second term is strictly positive, then $E[\hat{Q}(k)] \geq E[\hat{Q}(0)] \implies E[\hat{P}(k)] < E[\hat{P}(0)]$. This, in turn, implies that $\hat{q}(k) < \hat{q}(0)$. Finally, note that $\alpha > \frac{1}{2}$ implies $E[\hat{q}(\theta)] = \alpha^{\frac{\alpha}{1-\alpha}} \int \hat{P}(\theta)^{\frac{\alpha}{1-\alpha}} d\Phi(\theta) > (\alpha E[\hat{P}(0)])^{\frac{\alpha}{1-\alpha}} = \hat{Q}(0)$ by Jensen's Inequality establishing the assumed condition. The proof for $E[\hat{Q}(k)] < (=) E[\hat{Q}(0)]$ and $E[\hat{P}(k)] > (=) E[\hat{P}(0)]$ follows analogously. ■

If the rate at which marginal productivity decreases with output, itself falls with output, a firm adopting AI will supply a higher output on average than one without AI. Because the output of firms adopting AI and those not adopting AI are substitutes in demand, an expansion in the average output of firms adopting AI will contract the output of other firms. However, in equilibrium, the former effect outweighs the latter resulting in higher expected quantity and lower expected prices.

The next result of interest is with regard to the variance of observed market prices when k firms adopt AI. As shown by Nelson (1961), when a firm has access to a prediction of the underlying state variables that determine demand, their own elasticity of supply is inelastic. In the model here, $\hat{q}(\theta) = (\alpha \hat{P}(\theta))^{\frac{\alpha}{1-\alpha}}$ and, thus, the price elasticity of supply for an individual firm with AI is $\frac{dq}{dP}\frac{P}{q} = \frac{\alpha}{1-\alpha}(\alpha \hat{P}(\theta))^{\frac{2\alpha-1}{1-\alpha}} \frac{P}{(\alpha \hat{P}(\theta))^{\frac{\alpha}{1-\alpha}}} = \frac{\alpha}{1-\alpha}$. This is compared with the inelastic supply of a firm who does not adopt AI. Thus, as k firms now have elastic supply, the elasticity of aggregate supply is:

$$\epsilon_S(\theta) = k\left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\hat{P}(\theta)^{\frac{\alpha}{1-\alpha}}}{k\hat{P}(\theta)^{\frac{\alpha}{1-\alpha}} + (1-k)E[\hat{P}(\theta)]^{\frac{\alpha}{1-\alpha}}}\right)$$

Note that, ceteris paribus, elasticity is increasing in k and is 0 when $k = 0$.

Thus, we can establish:

Proposition 2 *When $k > 0$, the variance of $\hat{P}(\theta)$ is less than the variance of $\Phi(\theta)$.*

The proof trivially follows from the fact that θ is the intercept of linear market (inverse) demand. Thus, when supply is inelastic (i.e., as it is when $k = 0$), the variance of θ and $\hat{P}(\theta)$ are equal implying that when supply is elastic, the variance of $\hat{P}(\theta)$ falls.

Proposition 2 implies that the expected profit for firms adopting AI will be falling in k. As more firms adopt AI, the lower is the variance in prices and the lower are the benefits that can be achieved by matching firm output to price. The adoption decision, however, will be guided by the effect AI adoption has, if any, on the expected profits of firms who do not adopt AI. The following example with $\alpha = \frac{1}{2}$ shows that, as implied by Proposition 1, the profits of non-adoptees are unchanged as other firms adopt AI because expected market price is unchanged.

Example 2 *Returning to our previous example, note that the output choices of firms are:*

$$q = \tfrac{1}{2}(\rho(1 - (1-k)q - kq(1)) + (1-\rho)(2 - (1-k)q - kq(2))$$

$$q(\theta) = \tfrac{1}{2}(\theta - (1-k)q - kq(\theta))$$

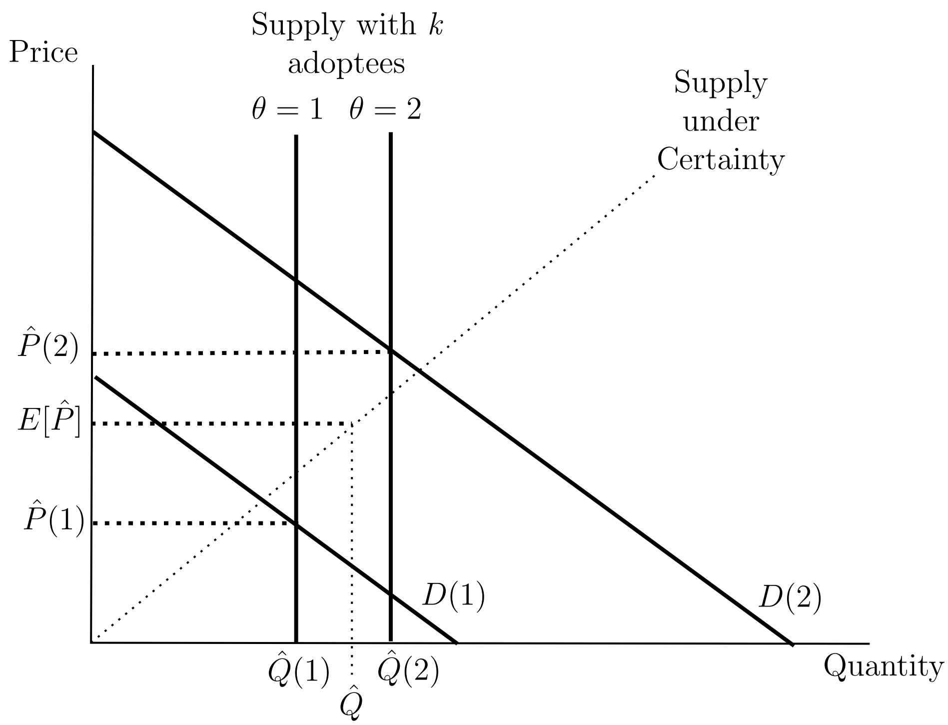

Figure 2: Market Equilibrium with AI Adoption by k firms

This implies: $\hat{q} = \frac{2-\rho}{3}$ *and* $\hat{q}(\theta) = \frac{2k+\rho(1-k)+\theta^2}{3k+6}$ *with* $\hat{P}(\theta) = \frac{4k+2(1-k)\rho+2\theta^2}{3k+6}$. *In Figure 2* $\hat{Q}(\theta) = (1-k)\hat{q} + k\hat{q}(\theta)$ *and* $\hat{Q} = E[\hat{Q}(\theta)]$. *Thus, the profit of a firm without AI is* $E[\hat{\pi}] = \frac{1}{9}(2-\rho)^2$, *the same as its profit when* $k = 0$ *while the expected profit of a firm with AI is:*

$$\rho(\hat{P}(1)\hat{q}(1) - (\hat{q}(1))^2) + (1-\rho)(\hat{P}(2)\hat{q}(2) - (\hat{q}(2))^2) = \frac{k(4+k)(2-\rho)^2 - \rho(5\rho+7)+16}{9(k+2)^2}$$

Critically, the derivative of this with respect to k *is* $-\frac{2\rho(1-\rho)}{(2+k)^3} < 0$. *Thus, the returns to adopting AI are falling in* k. *The reason for this is that, by Proposition 1, the expected market prices and quantities are invariant to* k *but the variance of prices is decreasing in* k. *That reduced variance, lowers the benefits of adjusting output to match demand and hence, the returns to adopting AI. Note, however, that as* $k \to 1$, *increment in expected profits from adopting AI becomes* $\frac{1}{9}\rho(1-\rho) > 0$.

2.3 Welfare: Firms

The above example shows that when $\alpha = \frac{1}{2}$, firms who do not adopt AI, earn the same profits regardless of k. What happens when $\alpha \neq \frac{1}{2}$? Proposition 1 shows that if $\alpha < \frac{1}{2}$, increased AI adoption will cause expected prices to rise. Thus, AI adoption confers an external benefit on non-adoptees. The opposite effect occurs if $\alpha > \frac{1}{2}$ with non-adoptees suffering a reduction in profit.

To understand the implications of this, suppose that it costs a firm, $c > 0$, if they want to adopt AI. Next let k^* be the share of firms adopting that maximise expected producer surplus. That is,

$$k^* \in \text{argmax}_k k(E[\hat{\pi}(\theta)] - c) + (1-k)E[\hat{\pi}]$$

This yields the following result:

Proposition 3 *Suppose that $E[\hat{\pi}(\theta)|0] - E[\hat{\pi}|0] \geq c$ and that each firm chooses independently whether to adopt AI or not. Then if $\alpha < (\geq)\frac{1}{2}$, $\hat{k} < (\geq)k^*$ firms adopt AI in equilibrium.*

Proof. Note that a firm will adopt AI if:

$$E[\hat{\pi}(\theta)|k] - E[\hat{\pi}|k] \geq c$$

Note that if $\alpha \geq \frac{1}{2}$ the externalities on non-adoptees are negative. Thus, if $E[\hat{\pi}(\theta)|0] - E[\hat{\pi}|0] \geq c$, then $E[\hat{\pi}(\theta)|k] - E[\hat{\pi}|k] > E[\hat{\pi}(\theta)|0] - E[\hat{\pi}|0]$ for all k and so all firms adopt AI. On the other hand, if $\alpha < \frac{1}{2}$ the externalities on non-adoptees are positive. Thus, the LHS of

$$E[\hat{\pi}(\theta)|\hat{k}] - E[\hat{\pi}|\hat{k}] = c$$

falls for $k > k^*$ and so $\hat{k} < k^*$. ■

Thus, producer surplus will be decreasing in k if $\alpha > \frac{1}{2}$ and non-decreasing otherwise.

2.4 Welfare: Consumers

Turning to consumers, in the absence of AI, while the quantity supplied by firms is the same regardless of demand, market prices fluctuate and so does consumer surplus. Given linear demand expected consumer surplus without AI ($E[CS_0]$) is given by:

$$E[CS_0] = \int \frac{1}{2}(\theta - \hat{P}(\theta))\hat{Q}d\Phi(\theta) = \frac{1}{2}\hat{Q}^2 = \frac{1}{2}\left(\alpha E[\hat{P}(\theta)]\right)^{\frac{2\alpha}{1-\alpha}}$$

If k firms adopt AI, then quantity supplied will adjust with demand. In particular, when θ is relatively high, quantity will expand while if θ is relatively low, it will contract. Thus, expected consumer surplus with k firms adopt AI ($E[CS_k]$) is given by:

$$E[CS_k] = \int \frac{1}{2}(\theta - \hat{P}(\theta))\hat{Q}(\theta)d\Phi(\theta) = \frac{1}{2}\int \hat{Q}(\theta)^2 d\Phi(\theta) = \frac{1}{2}\alpha^{\frac{2\alpha}{1-\alpha}} \int \hat{P}(\theta)^{\frac{2\alpha}{1-\alpha}} d\Phi(\theta)$$

Given this, observe that:

$$E[CS_k] \geq E[CS_0] \implies \int \hat{P}(\theta)^{\frac{2\alpha}{1-\alpha}} d\Phi(\theta) \geq \left(\int \hat{P}(\theta) d\Phi(\theta)\right)^{\frac{2\alpha}{1-\alpha}}$$

This will hold via Jensen's inequality if $\alpha \geq \frac{1}{3}$ and not otherwise.

Example 3 *Continuing the example, when $k = 0$, expected consumer surplus is $\frac{1}{18}(2-\rho)^2$. When there are k adoptees, consumer surplus is $\frac{k^2(\rho(5-8\rho)+4)+4k(\rho-2)^2+4(\rho-2)^2}{18(k+2)^2}$. Note that this is increasing in k; i.e., its derivative with respect to k is $\frac{2k\rho(1-\rho)}{(2+k)^3} > 0$.*

Intuitively, even though expected prices and quantities remain the same when $\alpha = \frac{1}{2}$, consumers benefit as firms use AI prediction to adjust their quantity supplied expanding it precisely when it is most valuable for consumers. Thus, for $\alpha \in \{\frac{1}{3}, \frac{1}{2}\}$, even though average prices rise and average quantity falls, consumers are still better off because they benefit more if an extra unit is produced in high θ outcomes than low θ outcomes. Over this range the interests of consumers and producers in adopting AI coincide. Outside this range, either consumers would prefer more AI (i.e., when $\alpha > \frac{1}{2}$) or producers would (i.e., when $\alpha < \frac{1}{3}$).

3 Labour Market Interactions

The analysis, thusfar, takes the wage rate, w, as exogenous to the firms in the industry. Here this assumption is relaxed and the wage rate is determined by equilibrium in the labour market. (The wage rate was also set equal to 1 but here it is re-introduced as a variable). If no firms adopt AI, their demand curve for labour is stable and would not cause variation in the realised wages in the labour market. However, if a sufficient number of firms adopt AI, their demand curves will be state dependent. Thus, as their output and hence, employment, choices are made *after* they receive a prediction of demand, this may impact on realised wages. In this situation, non-adoptees, if they know the number of firms who have adopted AI, may be able to infer θ from realised wages that they themselves must pay prior to production. Thus, in a manner reminiscent of Grossman & Stiglitz (1980), firms adopting AI are, by their actions, impacting upon observed price outcomes that provide an alternative signal to others of AI demand predictions they have received. The difference here is that it is not through single market asset trades that information is conveyed through prices but through vertical market effects; something that was previously explored by Gibbons et al. (2012) but here arises in a model not based on incomplete contracts.[5]

[5]Gibbons et al. (2012) were interested in a distinct set of research questions that arise in the context of firm boundaries under incomplete contracting and whether otherwise similar firms might choose distinct

In contrast to the assumption thusfar that labour was perfectly elastically supplied to the industry, here it is supposed that, in the short-run, the total supply of labour in the economy is fixed at $\widetilde{L}$. However, similarly to assumptions made by Gibbons et al. (2012), at any given time, that total supply is unknown and $\widetilde{L} \sim U[\underline{L}, \overline{L}]$. No firm observes $\widetilde{L}$. This uncertain supply implies that the wage in the labour market is not solely determined by the demand of firms in the industry.

Recall that, given k, labour demand is:

$$\hat{L} = k\left(\alpha\frac{P(\theta)}{w}\right)^{\frac{1}{1-\alpha}} + (1-k)\left(\alpha\frac{E[\hat{P}|w]}{w}\right)^{\frac{1}{1-\alpha}}$$

Setting $\hat{L} = \widetilde{L}$, this implies that:

$$\hat{w}(\theta, \widetilde{L}) = \left(\frac{1}{\widetilde{L}}\left(k(\alpha\hat{P}(\theta))^{\frac{1}{1-\alpha}} + (1-k)(\alpha E[\hat{P}|w])^{\frac{1}{1-\alpha}}\right)\right)^{1-\alpha}$$

Thus, when $k > 0$, the equilibrium wage potentially conveys information regarding θ with $\hat{w}$ non-decreasing in θ. As $\widetilde{L}$ is never observed, it is not necessarily possible for firms to infer θ from the realised wage. Nonetheless, as k rises, the signal that can be extracted from wages may become more precise. This will mean that firms who do not adopt AI may receive some of the benefits from AI indirectly from those who have adopted it and incurred c. Therefore, this positive externality from AI adoption means that a similar outcome to Proposition 3 arises and, from the perspective of producer surplus, too few firms adopt AI. Indeed, when $k = 1$, the wage becomes:

$$\hat{w}(\theta, \widetilde{L}) = \frac{\alpha\hat{P}(\theta))}{\widetilde{L}^{1-\alpha}}$$

Thus, even with high adoption, wages are not necessarily a perfect signal of demand.

Example 4 *Returning to our running example, note that, in the absence of any wage signalling effect, the realised wage is:*

$$\hat{w}(\theta, \widetilde{L}) = \frac{\sqrt{k^2(2-\rho)\left(2\theta^2 + 5\rho - 6\right) + k\left(\theta^2 + \rho\right)^2 + 4(2-\rho)^2}}{3(k+2)\sqrt{\widetilde{L}}}$$

Note that if $\bar{L} = \underline{L}$, then the lowest equilibrium wage consistent with $\theta = 1$, is $\hat{w}(1, \underline{L})$. Thus, if $\hat{w} > \hat{w}(1, \underline{L})$, this implies that θ cannot be 1. At the other extreme, if $\bar{L} = \overline{L}$, then the highest equilibrium wage consistent with $\theta = 2$, is $\hat{w}(2, \underline{L})$. Thus, if $\hat{w} < \hat{w}(2, \overline{L})$, this

organisational structures in equilibrium.

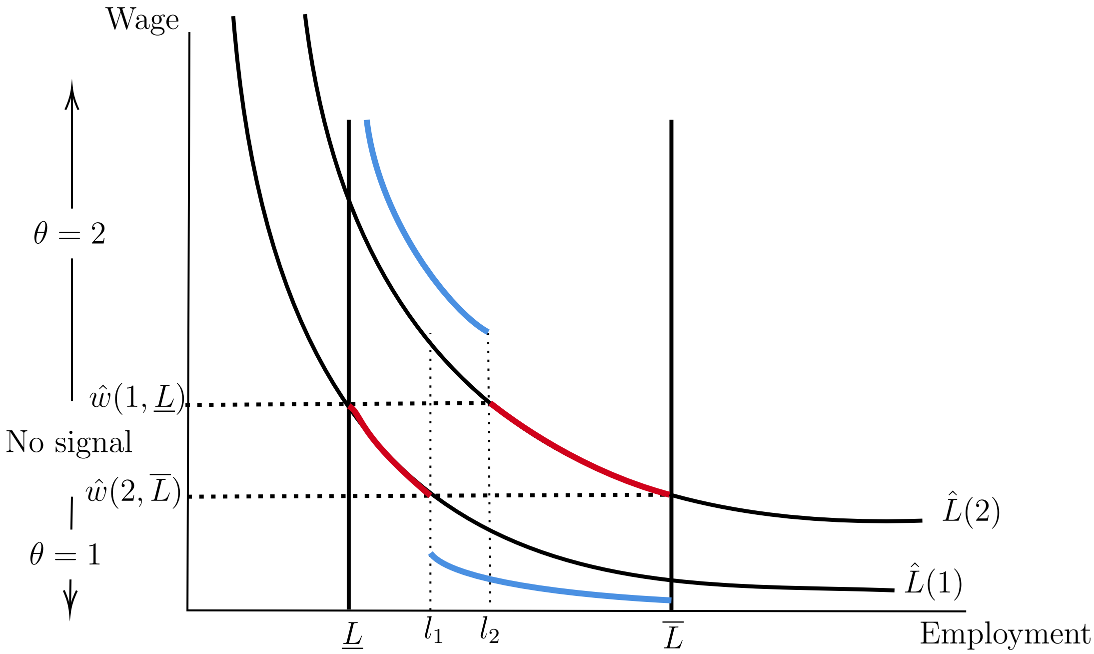

Figure 3: Labour Market Equilibrium with Adoption by k Firms

implies that θ cannot be 2. If, however, $\hat{w} \in [\hat{w}(2, \overline{L}), \hat{w}(1, \underline{L})]$ then, because $\hat{w}$ is continuous and strictly decreasing in $\widetilde{L}$, any realised wage in this range can be generated by two possible demand curves one associated with each value of θ. Thus, with probability ρ, $\theta = 1$ and with probability $1 - \rho$, $\theta = 2$ and the wage reveals no new information to a non-adopting firm (i.e., their posterior distribution remains identical to their prior).

If $\hat{w}(2, \overline{L}) > \hat{w}(1, \underline{L})$, there exist a range of w that will not be observed in equilibrium. In this case, a high(low) $\hat{w}$ is a signal that $\theta = 2(1)$. Note that there will be a range of $\hat{w}$ for which no signal of θ is sent if $\hat{w}(2, \overline{L}) \leq \hat{w}(1, \underline{L})$ or

$$\frac{\underline{L}}{\overline{L}} \geq \frac{k^2(\rho(8-5\rho)+4)+k(\rho+4)^2+4(2-\rho)^2}{k^2(2-\rho)(5\rho-4)+k(\rho+1)^2+4(2-\rho)^2}$$

Let $l_1 \equiv k\left(\frac{P(1)}{2\hat{w}(2,\overline{L})}\right)^2 + (1-k)\left(\frac{E[\hat{P}]}{2\hat{w}(2,\overline{L})}\right)^2$ and $l_2 \equiv k\left(\frac{P(2)}{2\hat{w}(1,\underline{L})}\right)^2 + (1-k)\left(\frac{E[\hat{P}]}{2\hat{w}(1,\underline{L})}\right)^2$. These are depicted in Figure 3. Note there that when $\widetilde{L} \in [l_1, l_2]$, then the clear signal means that the demand for labour increases when $\theta = 2$ and decreases when $\theta = 1$ for $k < 1$ (the blue lines). Outside these ranges labour demand is unchanged from before as no signal is sent to non-adoptees (the red lines). Given this, we can calculate the probability, $\rho\frac{\overline{L}-l_1}{\overline{L}-\underline{L}} + (1-\rho)\frac{l_2-\underline{L}}{\overline{L}-\underline{L}}$, that there is a clear signal of θ from $\hat{w}$:

$$\frac{3k(2k(\rho-2)-2\rho-5)\Big((\rho-1)\underline{L}\Big(k^2(\rho-2)(5\rho+2)-k(\rho+4)^2-4(\rho-2)^2\Big)+\rho\overline{L}\Big(-\Big(k^2(\rho-2)(5\rho-4)\Big)+k(\rho+1)^2+4(\rho-2)^2\Big)\Big)}{(k^2(\rho-2)(5\rho-4)-k(\rho+1)^2-4(\rho-2)^2)(k^2(\rho-2)(5\rho+2)-k(\rho+4)^2-4(\rho-2)^2)(\underline{L}-\overline{L})}$$

It is easy to show that this greater than 0 *and is increasing in* k*; that is, as more firms adopt the probability that a clear signal is generated by wage outcomes is higher. Note that as* $k \to 1$*, this probability becomes* $\frac{3(4\underline{L}(1-\rho)+\overline{L}\rho)}{4(\overline{L}-\underline{L})}$*. Thus, adoptions always generate a positive signal and if* $\frac{\underline{L}}{\overline{L}} \geq \frac{4-3\rho}{12(1-\rho)+4}$*, that probability is equal to* 1 *before everyone adopts.*

The example shows that AI adoption sometimes provides a perfect signal of product market demand while at other times provides no signal at all. The signal is provided when labour availability outcomes are extreme. For instance, if there is low labour availability but wages are very high, this is a signal that demand must also be relatively high. By contrast, if wages were low, this might indicate that either demand is low or that there is more labour available. As firms do not observe labour availability, the signal is muted. While the example provides a stark contrast between no signal and a perfect signal through wage rates, in reality, there will be imperfect signals. The important conclusion, however, is that the precision of the signals will rise as more firms are known to have adopted AI. Thus, separate from any price effect on non-adopting firms from AI adoption, there is a potential signaling effect through factor markets that reduces the marginal return to adopting AI as others adopt.

4 Long-Run Analysis

The short-run analysis shows how the adoption of AI prediction of demand allows firms in a competitive market to adjust the quantity of their variable input, labour, to match the forecast level of demand. As noted, while the returns to AI adoption vary with the degree of adoption in the market, using AI increases the expected profits of firms in the short-run. Thus, it is natural to examine what happens in the long-run.

It is assumed that the AI prediction of θ is sufficient to generate a short-run prediction only. Thus, only inputs variable in the short-run can be adjusted in response to that prediction. Inputs, such as capital, that only vary in the long-run, will, therefore, only adjust to the extent that these impact on expected short-run profits.

To begin, when capital is fixed at K_i in the short-run, for a firm without AI, its labour choice, resulting output and expected profit is:

$$\hat{L} = K_i \left(\alpha E[\hat{P}]\right)^{\frac{1}{1-\alpha}}$$

$$\hat{Q} = K_i \left(\alpha E[\hat{P}]\right)^{\frac{\alpha}{1-\alpha}}$$

$$E[\hat{\pi}|k=0] = K_i(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} E[\hat{P}|k=0]^{\frac{1}{1-\alpha}} - rK_i$$

where here we return to the notation where $w = 1$ as we do not explicitly consider the factor market signaling effect. Thus, it is easy to see that capital will flow into the industry to cause expected profits to fall to zero in equilibrium. This implies that:

$$E[\hat{P}|k=1] = \left(\frac{r}{(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}}}\right)^{1-\alpha}$$

Suppose alternatively that all firms in the industry find it profitable to adopt AI to maximise short-run profits. In this case,

$$E[\hat{\pi}|k=1] = K_i(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} \int \hat{P}(\theta|k=1)^{\frac{1}{1-\alpha}} d\Phi(\theta) - rK_i$$

Once again, these will fall to zero in equilibrium implying that:

$$\int \hat{P}(\theta|k=1)^{\frac{1}{1-\alpha}} d\Phi(\theta) = \left(\frac{r}{(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}}}\right)^{1-\alpha}$$

Let $\hat{K}_k$ be the long-run equilibrium level of capital in the industry given k. By Proposition 1, recall that, for a fixed level of capital, $E[\hat{P}|k=1,\hat{K}_k] < (\geq) E[\hat{P}|k=0,\hat{K}_k]$ if $\alpha > (\leq)\frac{1}{2}$. Note also that $\int \hat{P}(\theta|k=1,\hat{K}_k)^{\frac{1}{1-\alpha}} d\Phi(\theta) \geq E[\hat{P}|k=1,\hat{K}_k]$ by Jensen's Inequality. Thus, the following can be demonstrated:

Proposition 4 *The aggregate amount of capital $\hat{K}$ is higher when AI is adopted in the industry if $\alpha \leq \frac{1}{2}$. For any α, $E[\hat{P}|k=1,\hat{K}_1] \leq E[\hat{P}|k=0,\hat{K}_0]$.*

Proof. First note that $E[\hat{P}|k=0,\hat{K}_0] = \int \hat{P}(\theta|k=1,\hat{K}_1)^{\frac{1}{1-\alpha}} d\Phi(\theta) > E[\hat{P}|k=1,\hat{K}_1]$; proving the second part of the proposition. Second, if $\alpha \leq \frac{1}{2}$, $E[\hat{P}|k=1,\hat{K}_0] \geq E[\hat{P}|k=0,\hat{K}_0]$ which implies that $\hat{K}_1 > \hat{K}_0$. ■

This shows that, in the long-run, the availability and adoption of AI leads to lower expected prices and higher expected output in the industry. However, we can only unambiguously predict what happens to the aggregate amount of capital if α, the labour share, is less than $1-\alpha$, the capital share. Otherwise, it is possible that the aggregate level of capital could contract as AI is adopted. This is because when $\alpha > \frac{1}{2}$, it is possible that AI leads to lower levels of labour used on average and since these are complements with capital, can cause capital utilised to fall.

5 Conclusion

The tendency in the economic analysis of AI has been to equate AI adoption with automation. While there are applications of AI that are embodied in capital, AI is fundamentally an improvement in prediction technology and, therefore, its first order impacts will be to improve decision-making under uncertainty. When considering, therefore, the impact of its adoption of different classes of factors of production, AI adoption (say as measured by the employment of people with AI skills) will be a complement to inputs that are more variable in the short-run. AI prediction allows those inputs to be chosen to better respond to changes in variables and hence, increases their efficiency on average. If such variable inputs are in terms of labour and its employment, AI adoption will be primarily labour augmenting.

This paper has examined AI adoption from this perspective in a competitive market. Even here the adoption of AI has external effects on other firms. While most efficiency-enhancing innovations would involve adoption that reduced the profits of competing firms, here it is possible that AI adoption could increase the profits of competing firms. This may, in turn, limit the adoption of AI across such markets. It is shown that this arises, however, when variable inputs impacted on by AI are a smaller share of total inputs used by the firm. The broader AI's impact within a firm, the stronger will be the incentives to adopt and these will be driven by competition. This suggests that researchers will need to be careful in measuring the adoption and impact of AI in such markets.

There are many other issues that can be explored regarding the adoption of AI. One such avenue is to consider AI adoption by firms with market power – that is, who do not take market prices as given. This is explored by Gans (2022) who finds that the value of AI adoption differs upon whether it is informing pricing, output choices or both.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2020). *AI and jobs: Evidence from online vacancies*. Technical report, National Bureau of Economic Research.

Acemoglu, D. & Restrepo, P. (2018). Artificial intelligence, automation, and work. In *The economics of artificial intelligence: An agenda* (pp. 197–236). University of Chicago Press.

Acemoglu, D. & Restrepo, P. (2020a). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6), 2188–2244.

Acemoglu, D. & Restrepo, P. (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, D. & Restrepo, P. (2021). *Tasks, automation, and the rise in us wage inequality*. Technical report, National Bureau of Economic Research.

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019a). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31–50.

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019b). Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy*, 47, 1–6.

Brynjolfsson, E. (2022). The turing trap: The promise & peril of human-like artificial intelligence. *arXiv preprint arXiv:2201.04200*.

Brynjolfsson, E. & Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6370), 1530–1534.

Das, S., Steffen, S., Clarke, W., Reddy, P., Brynjolfsson, E., & Fleming, M. (2020). Learning occupational task-shares dynamics for the future of work. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 36–42).

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.

Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. In *AEA Papers and Proceedings*, volume 108 (pp. 54–57).

Gans, J. S. (2022). *AI Adoption in a Monopoly Market*. Technical report.

Gibbons, R., Holden, R., & Powell, M. (2012). Organization and information: Firms' governance choices in rational-expectations equilibrium. *The Quarterly Journal of Economics*, 127(4), 1813–1841.

Grossman, S. J. & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), 393–408.

Markoff, J. (2016). *Machines of loving grace: The quest for common ground between humans and robots*. HarperCollins Publishers.

Nelson, R. (1961). Uncertainty, prediction and competitive equilibrium. *Quarterly Journal of Economics*.

Pashigian, B. P. (1974). The accuracy of forecasts: A market equilibrium analysis. *The Quarterly Journal of Economics*, (pp. 86–97).

Susskind, D. (2020). *A world without work: Technology, automation and how we should respond*. Penguin UK.