NBER WORKING PAPER SERIES

AI ADOPTION IN A MONOPOLY MARKET

Joshua S. Gans

Working Paper 29995
http://www.nber.org/papers/w29995

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

AI Adoption in a Monopoly Market
Joshua S. Gans
NBER Working Paper No. 29995
April 2022
JEL No. D21,D81,O31

ABSTRACT

The adoption of artificial intelligence (AI) prediction of demand by a monopolist firm is examined. It is shown that, in the absence of AI prediction, firms face complex trade-offs in setting price and quantity ahead of demand that impact on the returns of AI adoption. Different industrial environments with differing flexibility of prices and/or quantity ex post, also impact on AI returns as does the time horizon of AI prediction. While AI has positive benefits for firms in terms of profitability, its impact on average price and quantity, as well as consumer welfare, is more nuanced and critically dependent on environmental characteristics.

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

While certain themes have been explored in economics regarding the adoption of AI including its role in labour replacement (Acemoglu & Restrepo (2018)) and in potentially facilitating collusion (Calvano et al. (2020)), there has been very little attention paid to how recent developments in AI will impact on the "meat and potatoes" operations of firms. That is, how will the adoption of AI change the price and quantity decisions of firms?

Usually, technological changes impact on those decisions through either process innovation (lowering the marginal costs of production and hence, reducing price and expanding quantity) or product innovation (improving demand and hence leading to price increases with ambiguous quantity implications). Overwhelmingly, the adoption of such innovations is seen as beneficial for both firms and consumers although it is possible to find exceptions (Bryan & Williams (2021)).

Some aspects of AI adoption do impact on firms like standard innovations. But, at its heart, recent AI developments are an advance in predictive statistics – allowing firms to generate and use information that was previously unavailable (see Agrawal et al. (2019)). For such innovations, the returns to adoption and impact on consumer welfare are not necessarily straightforward.

Here we explore one canonical class of predictions that (a) are valuable to most firms and (b) have clear implications for price and quantity decisions made by those firms. We look at predictions of firm demand. Through the gathering of larger datasets on consumers and more sophisticated multi-characteristic demand forecasting models using AI methods such as machine learning, in the future, firms may be able to predict demand precisely and further in advance of having to make key price and quantity decisions. This motivates us to work through the theory of how that improvement in information will impact on firm behaviour.

In this paper, the implications of moving from uncertain to certain demand are explored for a single monopoly firm.[1] The technical challenge in exploring this is not modelling price and quantity outcomes following AI adoption – those proceed along usual textbook lines – but modelling those choices prior to AI adoption. Specifically, as was noted many decades ago (Mills (1959)), when facing demand uncertainty, the price and quantity choices of a firm become challenging and do not collapse into a single dimension as they do textbook treatments. Moreover, different firms face different informational environments depending on the timing of decisions relative to the revelation of demand and also in terms of the time horizon of demand predictions. This gives rise to numerous cases and scenarios that must be parsed in order to develop a fulsome picture of the impact of AI adoption on firm choices.

[1] A companion paper looks at AI demand forecasting adoption by competitive firms (Gans (2022)).

This paper works through those cases in a way that allows an easy comparison of cases and characterisation of empirical and welfare implications.

Generally, it is found that AI adoption is beneficial to firms in terms of profit impacts but has ambiguous effects on price and quantity and hence, consumer welfare. This implies that there is no simple guide to the likely impact of AI adoption in a way that may influence regulatory considerations. This is not to say that some industry and other characteristics may guide regulators. It is to explain that just observing AI adoption by a monopolist is insufficient on its own.

The paper proceeds as follows. Section 2 shows the baseline results modelling the behaviour of the firm prior to AI adoption and comparing outcomes when AI is adopted in such a way that both price and quantity decisions can be made under certainty. Section 3 then turns to consider differing 'starting' environments for the firm depending on whether, prior to the adoption of AI, uncertainty was handled by targeting stock levels or by adjusting output to meet orders. These environments matter with AI adoption returns being higher for make-to-order rather than make-to-stock firms because the latter have more opportunities to exercise market power in AI's absence. Section 4 then looks at the time horizon of AI predictions and the returns to short-run decisions that inform one decision – price or quantity – but not the other. This allows a characterisation of when extending the time horizon of AI prediction involves increasing or decreasing returns to what we term AI "depth." A final section concludes.

2 Baseline Results

A monopoly firm with unit cost, $c \in (0, 1)$ and faces uncertain demand for a non-storable good given by:

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

We will assume that demand is linear with $Q = \theta - P$ where θ is an uncertain intercept that can take on one of two values, $\{1, 2\}$. The probability that $\theta = 1$ is $\rho < 1$. Some of the results below have antecedents with more general assumptions and these will be identified in the exposition that follows. As the focus here is on a comparison of various cases, using a simpler set of functional forms makes those comparisons more transparent.

2.1 Certainty benchmark (C)

It is useful to identify an outcome and performance benchmark. This will be the case when the firm has perfect information regarding θ prior to making any decision. In this case,

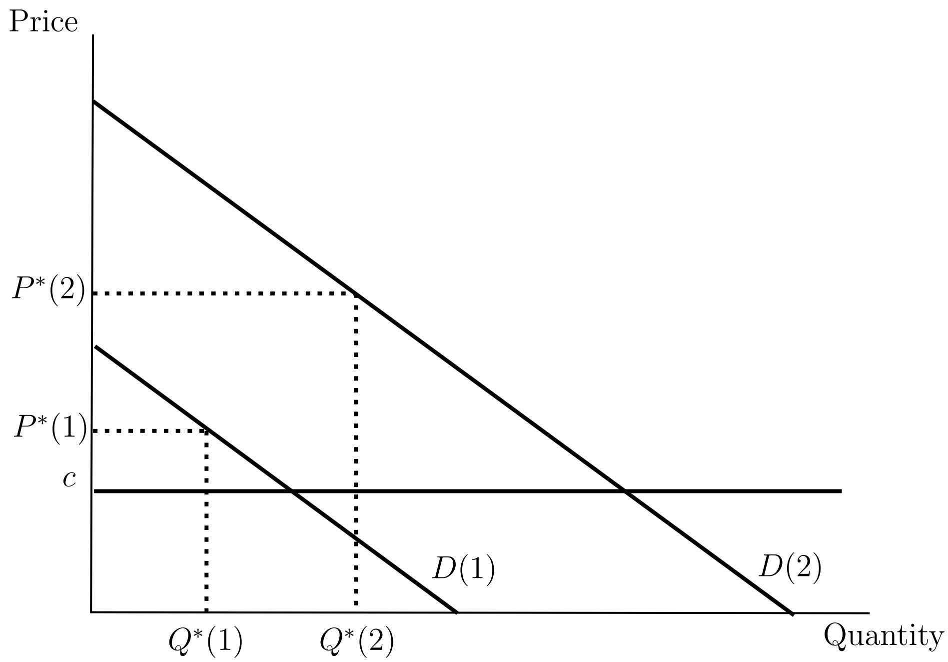

Figure 1: Price and Quantity Under Certainty (C)

$P^*(\theta) = \frac{1}{2}(\theta + c)$, $Q^*(\theta) = \frac{1}{2}(\theta - c)$ and $\pi^*(\theta) = (P - c)Q = \frac{1}{4}(\theta - c)^2$. Note also, that in terms of expected welfare, consumer, producer and total surplus are:

$$CS^* = \rho\tfrac{1}{2}(1 - P^*(1))^2 + (1 - \rho)\tfrac{1}{2}(2 - P^*(2))^2 = \tfrac{1}{8}\left((2c - 3)\rho + (2 - c)^2\right)$$

$$\bar{\pi}^* \equiv E[\pi^*(\theta)] = \rho(P^*(1)-c)(1-P^*(1))+(1-\rho)(P^*(2)-c)(2-P^*(2)) = \tfrac{1}{4}\left((2c - 3)\rho + (2 - c)^2\right)$$

$$TS^* = \tfrac{3}{8}\left((2c - 3)\rho + (2 - c)^2\right)$$

These outcomes are depicted in Figure 1. We will refer to this generically as case (C).

2.2 Price and quantity under uncertainty (U)

When θ is unknown, different outcomes arise depending on which decisions are made prior to θ being revealed. Here we consider the situation where uncertainty exists when *both* price and quantity are chosen.

In this case, the firm solves:

$$\max_{P,Q} P \min\{Q, E_\theta[D(P, \theta)]\} - cQ$$

This problem was first analysed by Mills (1959). Note that, by choosing $\{P^*(1), Q^*(1)\}$, the firm is always guaranteed $\pi^*(1)$. Thus, changing price and quantity will only be worthwhile if expected profits exceed this level.

Given the uncertainty, there are two types of errors that can be made.

1. (Surplus: Unsold inventory) $I(P,Q) = \rho(1-P-Q)$

2. (Shortage: Missed sales) $S(P,Q) = (1-\rho)(Q-(2-P))$

Based on these errors, expected profits become (using the decomposition of Lim (1980)):

$$PQ + PI(P,Q) - cQ = PE_\theta[D(P,\theta)] + PS(P,Q) - cQ$$

With general (rather than linear) demand, Mills (1959) showed that $\hat{P}_U \le E_\theta[P^*(\theta)]$, $\hat{Q}_U \ge E_\theta[Q^*(\theta)]$ and $\hat{\pi}_U \le E_\theta[\pi^*(\theta)]$. Thus, moving from uncertainty to certainty was profitable for the firm and, on average, raised price and lowered quantity.

The linear demand case gives us more insight into how uncertainty impacts on price and quantity decisions. We define $\Phi(c)$ as the level of ρ such that expected profits under uncertainty equal $\pi^*(1)$. We will show that $\Phi(c) = c + 2 - \sqrt{c(c+6)+1}$. Note that $\Phi(c)$ is decreasing in c with $\Phi(0) = 1$ and $\Phi(1) \approx 0.17$.

Proposition 1 *With linear demand,*

1. *If* $\rho > \Phi(c)$, $\hat{P}_U = P^*(1)$, $\hat{Q}_U = Q^*(1)$, $\hat{\pi}_U = \pi^*(1)$ *and* $CS_U \in [\frac{1}{2}Q^*(1)^2, \frac{1}{2}(\rho + (1-\rho)2 - P^*(1))Q^*(1)]$.

2. *If* $\rho \le \Phi(c)$, $\hat{P}_U = \rho P^*(1) + (1-\rho)P^*(2)$, $\hat{Q}_U = \frac{1}{2}(2+\rho-c)$, $\hat{\pi}_U = \pi^*(2) - \frac{1}{4}\rho(2c-\rho+4)$ *and* $CS_U = \frac{1}{2}\left(\rho(1-P^*(1)) + (1-\rho)(2-P^*(2))\right)\hat{Q}_U$.

Proof. In the linear demand case, expected profits collapses to:

$$\hat{\pi}_U = P(\rho(1-P) + (1-\rho)Q) - cQ$$

Note that this objective function is linear in Q. This implies that, but for a degenerate case, Q will be set as high or as low as possible depending on whether $(1-\rho)P > c$ or not. Note also that the lower (higher) level of Q will be optimally on the demand curve associated with $\theta = 1(2)$. Thus, expected profits are:

$$\pi = \begin{cases} (P-c)(1-P) & (1-\rho)P < c \\ (P-c)(2-P) - \rho P & (1-\rho)P > c \end{cases}$$

Thus, so long as $(1-\rho)P^*(1) < c$, profits are $\pi^*(1)$ with certainty. However, if $(1-\rho)P^*(1) > c$, $\hat{P}_U$ is the solution to $\max_P\{\rho(1-P)P + (1-\rho)(2-P)P - c(2-P)\}$ which gives $\frac{1}{2}(2-\rho+c)$.

It turns out this is equal to $\rho P^*(1)+(1-\rho)P^*(2)$ and $\hat{Q}_U = \frac{1}{2}(2+\rho-c) > Q^*(2)$. Intuitively, as $\hat{P}_U < P^*(2)$, it is optimal for the firm to produce a quantity in excess of $Q^*(2)$ that is the maximum amount it could sell at that price. Thus, we have:

$$\hat{\pi}_U = \begin{cases} \frac{1}{4}(1-c)^2 & \rho > \Phi(c) \\ \frac{1}{4}\left(c^2 - 2c(2+\rho) + (2-\rho)^2\right) & \rho < \Phi(c) \end{cases}$$

Note that if $(1-\rho)\hat{P}_U > c \implies (1-\rho)P^*(1) > c$ which accounts for the threshold $\Phi(c)$ on ρ (defining where the two profit cases hold with equality). Note also that when $\rho < \Phi(c)$, $I(\hat{P}_U, \hat{Q}_U) = -\rho$ and $S(\hat{P}_U, \hat{Q}_U) = 0$; that is, there are surpluses but no shortages, while when $\rho > \Phi(c)$, $I(\hat{P}_U, \hat{Q}_U) = 0$ and $S(\hat{P}_U, \hat{Q}_U) = -(1-\rho)$; that is, there are shortages but no surpluses.

When $\rho \leq \Phi(c)$, the consumer surplus calculations are straightforward as there are no shortages. However, shortages create an issue for the calculation of consumer surplus when $\rho > \Phi(c)$ and $\theta = 2$. In this case, when price is $P^*(1)$, quantity demanded is $2 - P^*(1) > Q^*(1)$. The open issue is which consumers are allocated the product if price is not performing a rationing function. One possibility is that the highest willingness to pay consumers obtain the product in which case consumer surplus is $\frac{1}{2}(2-P^*(1))Q^*(1)$. Another possibility is that the lowest willingness to pay customers obtain the product in which case consumer surplus is $\frac{1}{2}(2-(2-P^*(1)-Q^*(1))-P^*(1))Q^*(1) = \frac{1}{2}Q^*(1)^2$. Thus, expected consumer surplus is $\frac{1}{2}(\rho(1-P^*(1))+(1-\rho)Q^*(1))Q^*(1) = \frac{1}{2}Q^*(1)^2$. These extremes define the bounds in the proposition. ■

This outcome is depicted in Figure 2. The starting point to understand price and quantity choice under uncertainty is to note that the firm can always guarantee itself $\pi^*(1)$ by choosing $P^*(1)$ and $Q^*(1)$. Doing so, however, forgoes lost sales if demand is higher which happens with probability $1-\rho$ and, thus, the expected loss from not producing more while keeping price the same is $(1-\rho)P^*(1)$. When ρ is relatively high this is a worthwhile choice when $(1-\rho)P^*(1) \leq c$ because expanding output would result in losses. When, however, $(1-\rho)P^*(1) > c$, expanding output is potentially profitable. In this case, it is potentially worth producing as much as output as could be purchased at that price when demand is high. The cost to doing so is that when demand is low, there are unsold products. Because this is costly, it is profitable to raise price and lower quantity. However, it is not profitable to lower quantity all the way to $Q^*(2)$. This would result in no sales at all when demand is low even though products are being produced. Thus, instead, the firm chooses a higher quantity and lower price. That is, it is worth setting a lower price (at $\hat{P}_U$) so sales are made regardless of

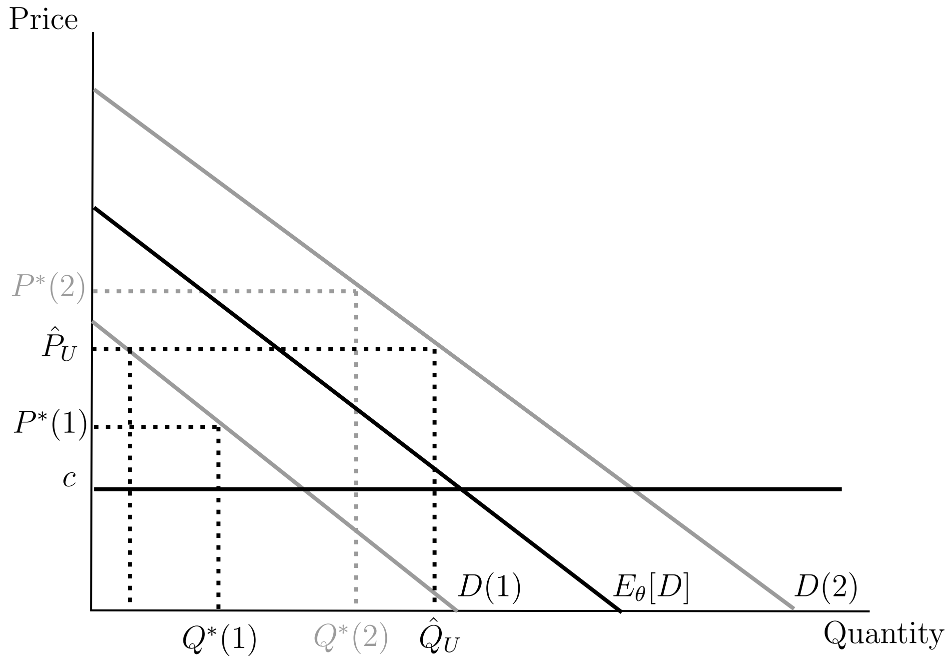

Figure 2: Price and Quantity Under Uncertainty (U)

demand and, because price is lower, it is worth expanding output to $\hat{Q}_U$.

With regard to welfare, when there is a (supply > demand) surplus, expected consumer surplus can be calculated in the standard way with the highest willingness to pay consumers being allocated the good as price is being used as a rationing device (this happens when ρ is relatively low and the firm chooses a high level of output). However, when there is a shortage (i.e., when ρ is relatively high), then price cannot be used to ration and so there is no clear approach as to which consumers are allocated the product and hence, what consumer surplus is. Thus, in the proposition, a range of consumer surplus is presented from the highest level (when high willingness to pay consumers are allocated the product) to the lowest level (when the lowest willingness to pay consumers are allocated the product).

2.3 The impact of AI adoption

AI offers a prediction of demand – or specifically θ. Suppose that (a) that prediction is perfect and reveals θ and (b) the prediction arrives before the firm chooses either price or quantity. As will be shown in Section 3 this second assumption may not be reasonable and incremental prediction is explored. For the moment, we consider the ability of AI adoption to more from complete uncertainty to complete certainty as a most favourable context for its adoption.

Given this, we can prove the following:

Proposition 2 *The returns to adopting AI are:*

$$\pi^* - \hat{\pi}_U = \begin{cases} (1-\rho)\frac{1}{4}(3-2c) & \rho > \Phi(c) \\ \rho\frac{1}{4}(1-\rho+4c) & \rho < \Phi(c) \end{cases}$$

If AI is adopted, expected consumer surplus falls while total surplus rises.

These are calculated from $\pi^* - \hat{\pi}_U$ and $CS^* - CS_U$. When ρ is relatively high, under uncertainty, the firm chooses a low price and quantity (generating $\pi^*(1)$) but with probability $1-\rho$ this forgoes the higher potential profits when $\theta = 2$. Thus, the return to AI is $(1-\rho)(\pi^*(2) - \pi^*(1))$. On the other hand, when ρ is relatively low, under uncertainty, the firm chooses to produce a relatively large amount. AI adoption allows them to economize on that cost and thus, the returns to AI are increasing in c.

From a consumer perspective, AI adoption reduces consumer surplus. When ρ is relatively high, AI allows price and quantity to rise when $\theta = 2$. But without AI, consumers still received the benefits of the increased demand and so their consumer surplus is higher than the certainty benchmark case. Thus, on average consumer surplus falls. When ρ is relatively low, the high production increased consumer surplus when demand was high and this outweighs any increase that might occur when demand is low under AI adoption. Once again, consumer surplus falls with AI adoption.

3 Incremental Decisions

The analysis of AI adoption thusfar has assumed that AI brings certainty to the firm from a starting point where neither price nor quantity decisions could be taken after the resolution of uncertainty. In reality, prior to AI, a firm may have one type of decision made under uncertainty while the other could be flexible and adjust when demand becomes known. For instance, a firm could solicit orders by committing to a price and then adjusting quantity depending on how many orders come in – this is what Milgrom & Roberts (1988) call a *make to order* environment (O).[2] Alternatively, a firm could engage in production and stock up on goods and then use prices afterwards to clear the market – this Milgrom & Roberts (1988) call a *make to stock* environment (S). In the former case, only price is set before demand uncertainty is realised while in the latter only quantity is set. In each case, therefore, what AI prediction brings is the ability to bring certainty to the incremental decision made under uncertainty in AI's absence. As will be seen, while the end point of adopting AI is the

[2] Lim (1980) refers to this as P-behaviour.

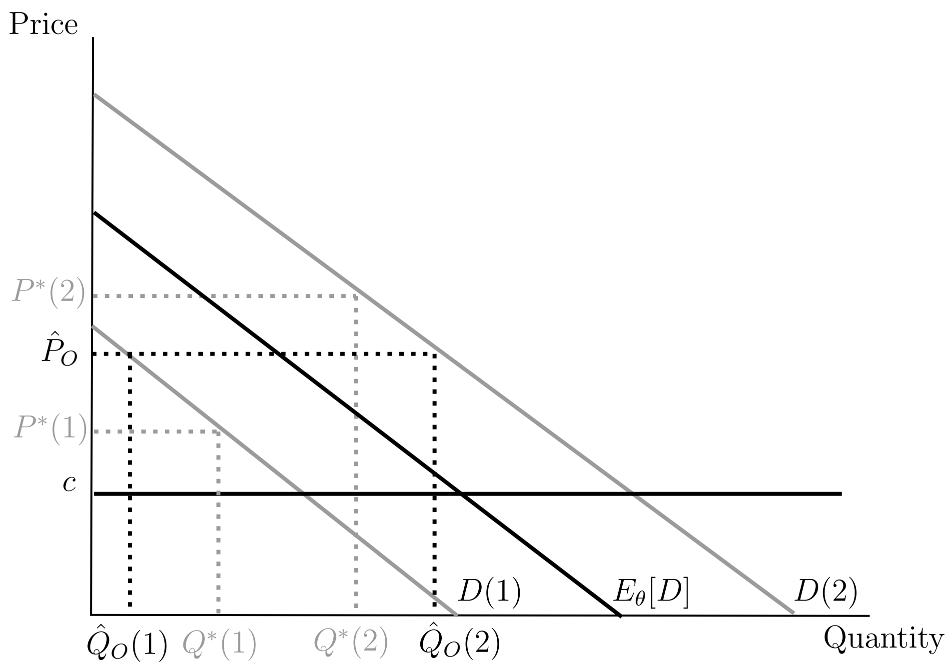

Figure 3: Price and Quantity Under Make to Order (O)

same – making both price and quantity decisions under uncertainty – the starting points are different. This impacts on the returns to AI as well as on its welfare implications.

3.1 Make to order (O)

Under make to order, price is chosen under uncertainty and consumers place orders based on that price. As the firm can then produce only for those orders it receives, quantity is adjusting ex post. Thus, there are no shortages or surpluses.

Given this, the firm solves:

$$\max_P (P - c)E_\theta[D(P, \theta)] = (P - c)(\rho + (1 - \rho)2 - P)$$

This yields $\hat{P}_O = \frac{1}{2}(2 - \rho + c)$, $E_\theta[\hat{Q}_P(\theta)] = \frac{1}{2}(2 - c - \rho)$ and $E_\theta[\hat{\pi}_P(\theta)] = \frac{1}{4}(2 - c - \rho)^2$. Note that $\hat{P}_O$ is the same as the $\hat{P}_U$ that arises when ρ is relatively low which itself equals $\rho P^*(1) + (1 - \rho)P^*(2)$. However, quantity can adjust. This means that while $\hat{Q}_O(2) = \frac{1}{2}(2 + \rho - c) = \hat{Q}_U$, $\hat{Q}_O(1) = \frac{1}{2}(\rho - c)$. These outcomes are depicted in Figure 3. Thus, $E_\theta[\hat{Q}_O(\theta)] \leq \hat{Q_U}$. In this case, $\hat{\pi}_O = \frac{1}{4}(2 - \rho - c)^2$. Of course, when ρ is high, it remains the case that $\hat{P}_O \geq \hat{P}_U$, $E_\theta[\hat{Q}_O(\theta)] \leq \hat{Q_U}$ but $\hat{\pi}_O > \hat{\pi}_U$ (as $\rho < 3 - 2c$ always).

AI adoption in a make to order environment involves a move to certainty (i.e., $O \to C$).

The return to AI adoption is:

$$\pi^* - \hat{\pi}_O = \tfrac{1}{4}\rho(1-\rho)$$

This is clearly a lower return than AI adopted in an environment of uncertainty.

With respect to consumer surplus, under make to order, expected consumer surplus is:

$$CS_O = \tfrac{1}{2}(\rho(1-\hat{P}_U)^2 + (1-\rho)(2-\hat{P}_U)^2) = \tfrac{1}{8}\left(2c\rho + (2-c)^2 - 3\rho^2\right)$$

Note that:

$$CS^* - CS_O = -\tfrac{3}{8}\rho(1-\rho)$$

Consumer surplus falls. This is because while AI adoption raises consumers surplus when there is low demand, it reduces it when there is high demand. Even when ρ is relatively high, the expected loss to high demand consumers outweighs that for low demand consumers. In terms of total surplus, there is an expected fall of $\frac{1}{8}\rho(1-\rho)$. Thus, adoption of AI, while profitable, reduces total welfare when coming from a make to order environment.

3.2 Make to stock (S)

When both price and quantity or simply price is chosen prior to the resolution of uncertainty, price is fixed and, thus, there is little room for the ex post exercise of market power. That isn't the case when a monopolist "makes to stock" (S). In that situation, upon the realisation of demand, the firm can still chose to restrict quantity below quantity produced in order to have an impact on price.

Before considering this, it is useful to suppose that ex post quantity reductions are not possible. That is, a firm must sell its stock. In this case, the price charged by the firm will adjust to match realised demand with available quantity. Thus, the firm solves:

$$\max_P (E_\theta[D^{-1}(Q,\theta)] - c)Q = (\rho + (1-\rho)2 - Q - c)Q$$

This yields $\hat{Q}_S = \hat{Q}_O$ and $E_\theta[\hat{P}_S] = \hat{P}_O$. In this case, $E_\theta[\hat{\pi}_S] = E_\theta[\hat{\pi}_O]$. As Lim (1980) shows, it is only when marginal costs are increasing or decreasing that differences arise.

In reality, there is no compulsion for a firm to sell rather than simply dispose of its stock. That is, it can make quantity restrictions ex post if it chooses. Thus, rather than passively adjust following the realisation of demand, a monopolist who has produced Q under uncertainty may choose to supply any $q \leq Q$ and influence the realised price, P. Working backwards, taking Q as given, if $\theta = 1$, the firm chooses $P_S(1)$ to maximise $P_S(1)(1-P_S(1)) - cQ$. This gives $\hat{P}_S(1) = \max\{\frac{1}{2}, 1-Q\}$ as Q and hence, total cost is taken as given. This

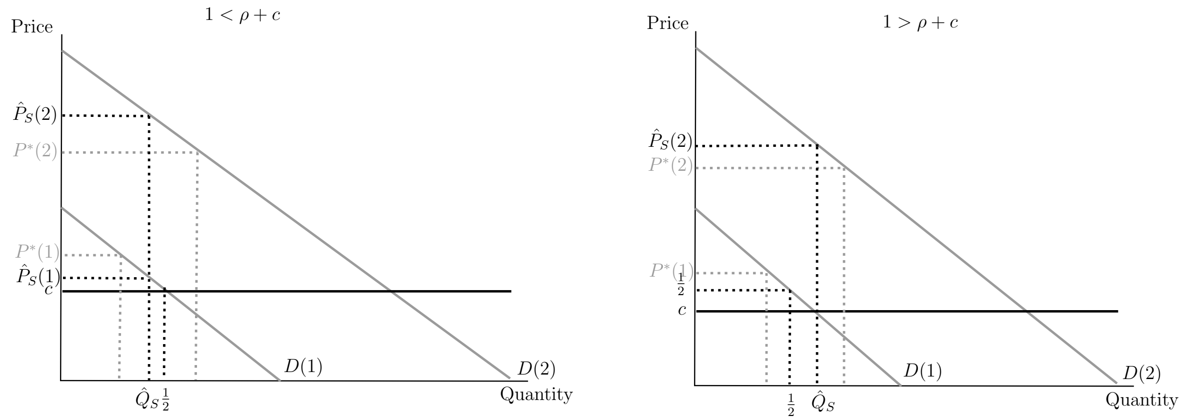

Figure 4: Price and Quantity Under Make to Stock (S)

yields profits of $\frac{1}{4} - cQ$ if $Q \geq \frac{1}{2}$ and $(1-Q-c)Q$ otherwise. On the other hand, if $\theta = 2$, the firm chooses $P_S(2)$ to maximise $P_S(2)(2-P_S(2)) - cQ$. This gives $\hat{P}_S(2) = \max\{1, 2-Q\}$ as Q and hence, costs are taken as given. This yields profits of $2-cQ$ if $Q \geq 1$ and $(2-Q-c)Q$ otherwise.

Given this, prior to the resolution of uncertainty, the firm chooses Q in anticipation of ex post pricing behaviour. If $Q < \frac{1-c}{2}$, then the firm will be unable to choose a price that allows for the lower bound of profits of $\frac{1}{4}(1-c)^2$ to be realised so it will set $Q \geq \frac{1-c}{2}$. Similarly a choice of $Q > \frac{2-c}{2}$ will be wasteful as $Q - 1$ will never be used. Thus, the quantity choice will lie in the range $[\frac{1-c}{2}, \frac{2-c}{2}]$. This means that Q will not exceed 1 so $P_S(2) = 2 - Q$

Thus, the firm chooses $Q \in [\frac{1-c}{2}, \frac{2-c}{2}]$ to maximise:

$$\rho \max\{\tfrac{1}{4}, (1-Q)Q\} + (1-\rho)(2-Q)Q - cQ$$

Note, first, that if $\hat{Q}_S < \frac{1}{2}$, then the firm will not want to restrict sales when $\theta = 1$. This is depicted in the left graph of Figure 4. Then the optimal (production) quantity is $\frac{2-\rho-c}{2}$ where $\hat{Q}_S < \frac{1}{2} \implies 1 < \rho + c$. In this case, expected profits are $\frac{1}{4}(2-\rho-c)^2$.

If this condition does not hold (i.e., $1 > \rho + c$), then the optimal quantity will be in the range $[\frac{1}{2}, \frac{2-c}{2}]$. Thus, when $\theta = 1$, sales will be restricted to $\frac{1}{2}$. This is depicted on the right graph of Figure 4. Then the optimal (production) quantity is $\frac{2(1-\rho)-c}{2(1-\rho)}$ which is greater than $\frac{1}{2}$ if $1 > \rho + c$. In this case, expected profits are $\frac{(4c-7)\rho+(2-c)^2}{4(1-\rho)}$.[3]

[3]Note that when $Q = \frac{1}{2}$, profits are $\frac{1}{4}(3 - 2(\rho + c))$. This is the optimal quantity if $1 = \rho + c$ implying that expected profits at this point are $\frac{1}{4}$. Thus, if $1 > (<)\rho + c$, $\hat{Q}_S > (<)\frac{1}{2}$.

To summarise, profits under make-to-stock are:

$$\hat{\pi}_S = \begin{cases} \frac{(4c-7)\rho+(2-c)^2}{4(1-\rho)} & 1 > \rho + c \\ \frac{1}{4}(2-\rho-c)^2 & 1 \leq \rho + c \end{cases}$$

Given this, it is easy to see that if $1 > \rho + c$, then profits exceed that case where sales passively adjust to output and also exceed the profits under make-to-order. This is because, in that case, ex post, the firm will restrict quantity made available for sale to lower than that available in order to maximise revenue. In effect, when $\theta = 1$, there is an opportunity to adjust prices to exercise market power even though this creates a surplus.

We are now in a position to analyse the impact of AI adoption. First of all, the returns to AI adoption are:

$$\pi^* - \hat{\pi}_S = \begin{cases} \frac{1}{4}\rho\frac{(2(1-\rho)-c)c}{1-\rho} & 1 > \rho + c \\ \frac{1}{4}\rho(1-\rho) & 1 \leq \rho + c \end{cases}$$

Note that if $1 \leq \rho + c$, the returns to AI are equal regardless of whether they are coming from make to order or make to stock. However, if $1 > \rho + c$, then $\pi^* - \hat{\pi}_S < \pi^* - \hat{\pi}_O$. This is because under make to stock, the firm is able to restrict quantity further when $\theta = 1$. This raises its profits in the absence of AI and hence, lowers the return to AI.

With respect to consumer surplus, under make to stock this is:

$$CS_S = \begin{cases} \frac{1}{2}(\rho(1-\frac{1}{2})\frac{1}{2} + (1-\rho)\hat{Q}_S(2)^2) = \frac{(4c-7)\rho+(2-c)^2+3\rho^2}{8(1-\rho)} & 1 > \rho + c \\ \frac{1}{2}(\rho + (1-\rho)2 - \hat{Q}_S)\hat{Q}_S = \frac{1}{8}((2-\rho)^2 - c^2) & 1 \leq \rho + c \end{cases}$$

While $CS^* < CS_S$ when $1 < \rho + c$, interestingly, AI adoption improves consumer surplus when $1 > \rho + c$. This is because while it raises prices and lowers quantity when $\theta = 1$, it results enough of an increase in both quantity and prices when $\theta = 2$ which outweighs the low demand effect. Thus, it is precisely when a firm would otherwise dispose of stock to exercise market power, that AI adoption unambiguously increases welfare.[4]

3.3 Comparisons

Make-to-stock and make-to-order represent two starting points for a firm without AI prediction of demand. It is instructive to compare the returns to AI adoption in each case in order to generate an empirical implication as to which type of firm/industry is more likely to see AI adoption.

[4]Note that, if $1 < \rho + c$, $CS_O > CS_S \implies (\rho - c)(2(1-\rho) - c) > 0$. Thus, there is no simple ranking of consumer surplus from these options.

The following proposition compares the returns to AI adoption in each case.

Proposition 3 *The returns to AI adoption are higher for make-to-order firms than make-to-stock firms if $\rho < 1 - c$ and are the same otherwise.*

The proof follows from that fact that $\pi^* - \hat{\pi}_S < \pi^* - \hat{\pi}_O \implies \frac{1}{4}\rho\frac{(2(1-\rho)-c)c}{1-\rho} > \frac{1}{4}\rho(1-\rho)$ when $1 > \rho + c$. Intuitively, a firm who is using make-to-stock has higher expected profits than one using make-to-order when c is relatively low because the cost of have surplus production (i.e., in practice, holding inventory) is itself relatively low. Thus, the firm benefits from the option of using market power to keep prices high should demand turn out to be low. Adopting AI eliminates uncertainty for both firm types. However, by eliminating the reason to have an option to exercise market power, consumer surplus actually increases when AI is adopted by make-to-stock firms when c is low.

3.4 Comparison with competition

In a companion paper (Gans (2022)), AI adoption is a competitive market is examined. In that situation, firms choose their quantity under uncertainty and price adjusts to equate supply and demand. Thus, AI adoption here is comparable to "make-to-stock" for a monopoly firm.

What happens in the environment here if there are many firms choosing quantity under uncertainty taking price as given? As Gans (2022) shows, firms choose quantity such that expected price is equal to marginal cost. That is, $\rho\hat{P}(1)+(1-\rho)\hat{P}(2) = c$ where $\hat{P}(1) = 1-Q$ and $\hat{P}(2) = 2 - Q$ where Q is the aggregate quantity chosen by all firms. It can be easily seen that $\hat{Q} = 2 - c - \rho$. Expected consumer surplus is $\frac{1}{2}(2 - c - \rho)^2$.

If all firms adopt AI, then quantity adjusts with demand so that $1 - \hat{Q}(1) = c$ and $2 - \hat{Q}(2) = c$. Thus, expected quantity is $\rho + (1 - \rho)2 - c$ and expected price is c. Expected consumer surplus is $\frac{1}{2}\left((2c - 3)\rho + (c - 2)^2\right)$. Thus, AI adoption increases consumer surplus by $\frac{1}{2}\rho(1 - \rho)$.

This implies that expected consumer surplus increases by more as a result of AI adoption under competition than monopoly. This occurs even where $\rho < 1 - c$ where AI adoption increases consumer surplus moving from $U \to S$. That said, in a competitive market, as shown by Gans (2022), there are externalities between firms in terms of the incentives to adopt AI that may mitigate its adoption in competitive markets whereas for a monopolist adopting AI is always profitable.

4 Short-Run Predictions

In each of the cases considered thusfar, AI prediction allows the firm to resolve uncertainty prior to choosing *both* price and quantity. The difference was in the starting point regarding whether both price and quantity or one or the other was chosen under uncertainty.

In some situations, however, what AI can achieve may be more limited. Rather than providing a long-run prediction that allows both price and quantity decisions to be made under certainty, it may be that AI prediction is short-run – sufficient to allow one type of decision to be made under certainty. For instance, AI may permit firms to choose a price based upon a demand prediction but not output. In this case, AI adoption involves $U \to O$, $S \to C$ or no change as the case may be. Alternatively, AI may permit firms to choose a quantity based on a demand prediction even if price is committed to earlier. AI adoption then involves $U \to S$, $O \to C$ or no change. Here we examine the returns and impact of AI adoption for short-run predictions.

4.1 Price under certainty

Suppose, first, that AI adoption permits price to be chosen under certainty about demand (if it isn't already). We had previously analysed the $S \to C$ case in Section 3.2 above, thus, we state the returns to AI adoption and impact on consumer surplus for the $U \to O$ case.

$$\hat{\pi}_O - \hat{\pi}_U = \begin{cases} \frac{1}{4}(1-\rho)(3-2c-\rho) & \rho > \Phi(c) \\ \rho c & \rho < \Phi(c) \end{cases}$$

$$CS_O - CS_U = \begin{cases} \left[\frac{1}{8}(1-\rho)(3\rho+3-2c), \frac{1}{8}(\rho(2-3\rho)+1)\right] & \rho > \Phi(c) \\ \frac{1}{4}\rho(c-\rho) & \rho < \Phi(c) \end{cases}$$

Recall that under U, when there is a shortage, there is a range of potential consumer surplus and this range is translated into a range of differentials. Nonetheless, in each case, the increment to consumer surplus is positive. Thus, as in the case of $S \to C$ when $1 > \rho + c$, AI adoption unambiguously increases welfare.

4.2 Quantity under certainty

Now suppose that AI adoption permits quantity to be chosen under uncertainty about demand (if it isn't already). We had previously analysed the $O \to C$ case in Section 3.1 above.

Here the relevant new case of interest is $U \to S$.

$$\hat{\pi}_S - \hat{\pi}_U = \begin{cases} \frac{1}{4}(1-\rho)(3-2c-\rho) & \rho > 1-c, \Phi(c) \\ \rho c & 1-c < \rho < \Phi(c) \\ \frac{1}{4}\left(\frac{c^2\rho}{1-\rho} - 2c + 3(1-\rho)\right) & 1-c > \rho > \Phi(c) \\ \frac{\rho(1-\rho+c)^2}{4(1-\rho)} & \rho < 1-c, \Phi(c) \end{cases}$$

$$CS_S - CS_U = \begin{cases} \left[\frac{2c(1-c)+\rho^2-4\rho+3}{8}, \frac{2c(1-c)+(2c+1-\rho)(1-\rho)}{8}\right] & \rho > 1-c, \Phi(c) \\ \frac{1}{4}(c-\rho)(2-c-\rho) & 1-c < \rho < \Phi(c) \\ \left[\frac{c^2\rho+(1-\rho)(3(1-\rho)-2c)}{8(1-\rho)}, \frac{1-\rho(c(2(1-\rho)-c)-2+\rho)}{8(1-\rho)}\right] & 1-c > \rho > \Phi(c) \\ \frac{\rho\left(c^2-\rho^2+4\rho-3\right)}{8(1-\rho)} & \rho < 1-c, \Phi(c) \end{cases}$$

All the ones with the range of CS are positive (that is, where under uncertainty there would be rationing). The last one is always negative while the second is negative if $\rho > c$. In these cases, under uncertainty, there is a large quantity supplied and so, when adopting AI, this dramatically reduces the quantity available and hence, reduces consumer surplus.

4.3 Returns to AI "Depth"

Short-run predictions are assumed to enable the resolution of uncertainty for just one decision – either quantity or price. Thus, they facilitate changes impacting the quantity decision – $U \to O$ and $S \to C$ – or the price decision – $U \to S$ and $O \to C$. Note that, for the same reason that $\pi^* - \hat{\pi}_O > \pi^* - \hat{\pi}_S$ (when $1 > \rho + c$), $\hat{\pi}_S - \hat{\pi}_U > \hat{\pi}_O - \hat{\pi}_U$ (namely, that $\hat{\pi}_S > \hat{\pi}_O$). The question we want to ask here is: what is the shape of the returns to AI "depth"; that is, investing beyond a short-run prediction to a longer-run prediction?

Formally, this question asks whether $\pi^* - \hat{\pi}_O > \hat{\pi}_O - \hat{\pi}_U$ and/or $\pi^* - \hat{\pi}_S > \hat{\pi}_S - \hat{\pi}_U$? An affirmative answer would indicate that there are increasing returns to AI "depth".

Proposition 4 *There are increasing returns to AI "depth" through S (i.e.,* $\pi^* - \hat{\pi}_S > \hat{\pi}_S - \hat{\pi}_U$*) only if* $\rho > \frac{3}{2} - c$ *and decreasing returns otherwise. There are increasing returns to AI "depth" through O (i.e.,* $\pi^* - \hat{\pi}_O > \hat{\pi}_O - \hat{\pi}_U$*) only if* $\rho > \frac{3}{2} - c$ *and decreasing returns otherwise.*

Proof. First, consider the S route. Note that:

$$\pi^* - \hat{\pi}_S - (\hat{\pi}_S - \hat{\pi}_U) = \begin{cases} \frac{1}{4}(1-\rho)(2c+2\rho-3) & \rho > 1-c, \Phi(c) \\ \frac{1}{4}\rho(1-\rho-4c) & 1-c < \rho < \Phi(c) \\ \frac{-2c^2\rho+2c\left(1-\rho^2\right)-3(1-\rho)^2}{4(1-\rho)} & 1-c > \rho > \Phi(c) \\ \frac{\rho\left(-2c^2-(1-\rho)^2\right)}{4(1-\rho)} & \rho < 1-c, \Phi(c) \end{cases}$$

All these are negative but the first case which can be positive if $\rho > \frac{3}{2} - c$.

Second, consider the O route. Note that:

$$\pi^* - \hat{\pi}_O - (\hat{\pi}_O - \hat{\pi}_U) = \begin{cases} \frac{1}{4}(1-\rho)(2c+2\rho-3) & \rho > \Phi(c) \\ \frac{1}{4}\rho(1-\rho-4c) & \rho < \Phi(c) \end{cases}$$

These are negative except for the first case which can be positive if $\rho > \frac{3}{2} - c$. ■

This proposition says that, generically, there are decreasing returns to AI "depth" except where ρ is sufficiently high. A high ρ implies that low demand occurs relatively frequently. In this case, moving one decision out of two to be made when there is information about demand has relatively little impact on eliminating the costs of uncertainty and, thus, there is value to AI that will allow both decisions to be made under certainty in a coordinated fashion. That is, the firm acts "as if" demand is always low in these cases unless both price and quantity can adjust to higher demand.[5]

5 Conclusion

This paper has demonstrated that while AI presents an opportunity for businesses, its impact on key choices that might be of regulatory interest – i.e., prices and quantities – is far from straightforward. Moreover, these results were presented in a somewhat idealised environment with a firm facing a simple form of uncertainty, having constant marginal cost and AI itself offering perfect predictions. In the real world, numerous complexities abound that should give economists pause when trying to draw general conclusions regarding the impact of AI

[5] Another way of framing this issue is that there may be different divisions in a firm responsible for the price and quantity decisions. It may be that the division in charge of pricing (say, marketing) receives information from the AI on demand while the division in charge of quantity (say, operations) does not. In the absence of communication between the two, AI adoption would involve $U \to O$ while by investing in communication, the decisions could be coordinated and result in a move from $U \to C$. The proposition shows that such coordination will be important for AI adoption when $\rho > \frac{3}{2} - c$. Thus, it provides a specific example of the system issues examined by Agrawal et al. (2021).

adoption.

References

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.

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

Agrawal, A. K., Gans, J. S., & Goldfarb, A. (2021). *AI Adoption and System-Wide Change*. Technical report, National Bureau of Economic Research.

Bryan, K. A. & Williams, H. L. (2021). Innovation: market failures and public policies. In *Handbook of Industrial Organization*, volume 5 (pp. 281–388). Elsevier.

Calvano, E., Calzolari, G., Denicolo, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–97.

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

Lim, C. (1980). The ranking of behavioral modes of the firm facing uncertain demand. *The American Economic Review*, 70(1), 217–224.

Milgrom, P. & Roberts, J. (1988). Communication and inventory as substitutes in organizing production. *The Scandinavian Journal of Economics*, (pp. 275–289).

Mills, E. S. (1959). Uncertainty and price theory. *The Quarterly Journal of Economics*, 73(1), 116–130.