NBER WORKING PAPER SERIES

PREDICTION MACHINES, INSURANCE, AND PROTECTION: AN ALTERNATIVE PERSPECTIVE ON AI'S ROLE IN PRODUCTION

Ajay K. Agrawal
Joshua S. Gans
Avi Goldfarb

Working Paper 30177
http://www.nber.org/papers/w30177

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

Thanks to the Brookings Foundation for funding. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Prediction Machines, Insurance, and Protection: An Alternative Perspective on AI's Role in Production
Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb
NBER Working Paper No. 30177
June 2022
JEL No. D81,O32

ABSTRACT

Recent advances in AI represent improvements in prediction. We examine how decision-making and risk management strategies change when prediction improves. The adoption of AI may cause substitution away from risk management activities used when rules are applied (rules require always taking the same action), instead allowing for decision-making (choosing actions based on the predicted state). We provide a formal model evaluating the impact of AI and how risk management, stakes, and inter-related tasks affect AI adoption. The broad conclusion is that AI adoption can be stymied by existing processes designed to address uncertainty. In particular, many processes are designed to enable coordinated decision-making among different actors in an organization. AI can make coordination even more challenging. However, when the cost of changing such processes falls, then the returns from AI adoption increase.

Ajay K. Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
and NBER
ajay.agrawal@rotman.utoronto.ca

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

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
and NBER
agoldfarb@rotman.utoronto.ca

1 Introduction

Artificial intelligence (AI) advances have been coming at a rapid pace for the past decade. This has been accompanied by increasing speculation about the potential impact of AI. What are AI's complements and substitutes? The research to date has primarily focused on the role AI will play in substituting for humans in the performance of specific occupational tasks. However, recent advances in machine learning constitute advances in the statistics of prediction and, hence, are substitutes for human prediction (Agrawal et al. (2018b)). In this paper, we consider the role of prediction in decision-making under uncertainty and evaluate its potential to substitute for other ways risk in decisions is managed. Relative to the dominant view of AI as a technology for labor substitution for individual tasks, the distinct perspective provided here focuses on precisely where the impact of AI in the economy is likely to be. Taking this view focuses attention on the factors employed in risk management strategies.

Seen as prediction, a key question for understanding the impact of AI relates to how it changes the role of decisions in organizations. In particular, we examine how AI might enable different types of decision-making. Having better signals about the environment leads to the explicit use of those signals when actions are chosen. In the absence of those signals, production processes are organized around rules. Rules are actions made that do not adjust to their environment as those choosing the action do not have signals that would cause such adjustment. Put differently, without predictions, it might be best to always to the same thing (to follow a rule). Prediction enables choices (decisions).

There are two consequences of a change from rules to decisions. First, rules do not change as new information arises, implying that there are efficiency consequences that depend on the realization of that uncertainty. This causes organizations to set up alternative activities designed to support rules. These alternatives, which Ehrlich & Becker (1972) labeled 'insurance' (that mitigates losses) and 'protection' (that reduces the probability of bad outcomes), reduce the negative consequences of a bad outcome. The focus of this paper is to examine how decision-making and risk management strategies change when prediction improves, building on the framework in Agrawal et al. (2019).

If prediction substitutes away from rules towards decisions, the need for those activities is reduced. Second, rules create a certain degree of reliability in organizations where there are a number of interrelated decisions. If everyone else is always doing the same thing, then coordination is straightforward. If prediction triggers a shift to decisions, this reduces reliability and may reduce the productivity of complementary tasks. It is this potential that caused Bresnahan (2020) to conclude that AI adoption would only occur where such interrelationships did not exist – that is, when AI is adopted for a modular task. By contrast, we argue that this potential creates new opportunities for organizational or system change that itself requires new activities designed to more explicitly coordinate tasks to avoid the negative consequences of reduced reliability.[1]

In contrast, much of the existing AI literature focuses on the impact of AI on labor with an emphasis on automation of human tasks by thinking machines (e.g. Felten et al. (2018), Acemoglu & Restrepo (2019), Rock (2019), Frank et al. (2019), Das et al. (2020), Acemoglu (2021)). The fundamental assumption driving these approaches is that AI improvements are a substitute for cognition at a task level. This view is related to a conception of AI that is quite common in popular culture and also in the aspirations of computer scientists pushing for what is termed "artificial general intelligence" (AGI). The problem is that *this is not the technological advance that computer science over the last decade has achieved.* Instead, the advances that we have seen are in machine learning which is more properly characterized as an advance, albeit a significant one, in statistics; in particular, in statistics to generate predictions (Agrawal et al. (2018b)).

Our focus on AI as prediction suggests that AI's impact will be primarily in organizations structured around uncertainty. This viewpoint is a significant challenge to the current presumption in economics that AI will, at a first order, spur capital to labor substitution and, indeed, to the notion that AI adoption is human cognition substitution. Instead, it is more consistent with Bresnahan (2020), who emphasizes that AI adoption will change systems. Our results suggest that AI will substitute for insurance and protection. There is no reason to suppose that these insurance and protection activities involve a lower capital intensity than those in the economy. Thus, their removal may reduce the average capital intensity of production. Moreover, to the extent that AI prediction opens up the explicit decisions in tasks rather than rules, there is an increased use of human cognition in the form of the application of judgment (Agrawal et al. (2018a)).

More generally, the empirical implications of our framework are not about the identification of tasks that AI can do, nor about the types of workers that AI will replace. Instead, our framework provides insight into the types of organizations that will benefit from AI, and the challenges to AI adoption more generally. Our framework suggests that AI adoption will be more straightforward when uncertainty constrains what an organization can do, and when multiple decisions can be easily coordinated through communication. When communication between decision-makers is difficult, or when uncertainty plays little obvious role in how an organization operates, AI adoption will be limited.

The paper proceeds as follows. In Section 2, we discuss examples of insurance (reducing

[1] This is explored in more detail in Agrawal et al. (2021).

the consequences of a bad state should it occur) and protection (reducing the chance of a bad state occurring) in dealing with uncertainty, and the use of rules by many organizations to reduce that uncertainty. This discussion will hint at a role for better prediction to reduce the need to depend on rules. After this motivating discussion, in Section 3, we introduce a formal model to explore the conditions under which is it beneficial to switch from using a rule to using a decision. We define protection and insurance, and demonstrate that when prediction (through AI) is good enough, it reduces the need for protection and insurance. This enables organizations to move from rules where they do the same action regardless of the state to decisions. However, when the stakes of the decision are high, specifically when the payoffs are asymmetric and getting the decision wrong can be particularly consequential, we show that rules maintain their appeal and the returns to AI adoption are lower. In Section 4, we explore how insurance, protection, and prediction interact when there are multiple decisions within an organization. The interaction of decisions leads to a phenomenon we call the AI Bullwhip Effect. The need to coordinate can mean that AI complements insurance, though not protection. We conclude with a discussion of the implications of the model for understanding the impact of AI more broadly and offer some directions for future research.

2 Risk Management Activities

Our formal model will show that what AI prediction has the potential to substitute for are other ways in which the risk associated with actions taken under uncertainty are managed. Before moving to the formalities, it is useful to discuss briefly what we mean by risk management.

To adopt AI, you must have something to predict. Specifically, there has to be something about the world that is uncertain and that you believe will impact your business. In some cases, uncertainty is obvious. Credit card companies know they don't know which transactions are fraudulent. Amazon knows they don't know what their customers want to buy at any given time. These happen to be industries where the uncertainty is there for all to see.

But what if that isn't the case? Uncertainty is costly and there are many things that we do to protect ourselves from that uncertainty. If we have done this successfully, we might no longer be aware that uncertainty was there, let alone that it could be embraced more directly and overcome using AI prediction. Consider airport terminals. Many of these are currently designed to be a destination in and of themselves. The reason is that it is expected that many people will spend time waiting in airports. Hence, it is an opportunity to provide commercial activities that can be consumed while waiting – e.g., shopping but in some cases, sports and entertainment facilities. If that goal is achieved, people will be less concerned about the time they spend at airports and hence, it will be less obvious that changes – such as AI prediction to optimize airport arrival times or security line transit – will be valued. Contrast this with private terminals that are spartan precisely because, in their case, travellers do not have to adhere to a schedule and instead, flexibility allows them to minimize any waiting time.

2.1 Examples of insurance

Retail inventory: To see how risk management plays a role in concealing uncertainty, consider the canonical task of acquiring a stock of products prior to knowing fully what demand will be. While it is possible to match supply with demand if your customers are willing to wait for orders to be filled, when customers have a preference for fast fulfillment, goods must be produced prior to precisely knowing what demand will be. In this case, while a retailer may luck out and order the exact amount demanded, invariably there are two types of mismatches that can arise – a supply shortage or a surplus.

In choosing how much to stock, a seller must balance these two mismatches. A shortage means that you have forgone sales. A surplus means that you have unsold inventory. However, only one sort of mismatch provides options for insurance (for reducing the consequences of being in the worse state). If you have a shortage, those customers may never return. On the other hand, if you have a surplus, you have options. You might discount your product to encourage more sales or you might be able to re-sell them in other markets or store them to sell later. Only if these options do not exist, as they might not for perishable goods, do you have to dispose of the goods and book the loss. Thus, precisely because insurance is a strong option, many sellers tend to order more than they might expect to sell.

Perhaps nowhere do we see this attention to an insurance type strategy more clearly than in fashion. High quality garments take time to produce. That means that designers and manufacturers must forecast demand a season or more in advance. But fashion purchases are complex. They are not simply driven by the need for clothing. Most consumers already have clothes. Instead, they are also driven by a preference to send social messages. If designers anticipate those social messages correctly, demand will be high. If they do not do so – say, by picking colors that turn out to be unappealing – then demand may be very low. In 2018, Burberry threw away stock worth $36 million. What's more, that value was six times more than they had to deal with in 2013. Rare are fashion items that sell out. Much more common is a large volume of unsold stock at the end of a season.

Fashion waste has led to a raft of insurance-like choices. There are perennial end of season discount sales and a host of outlet malls to unload excess stock. Some retailers ship unsold stock to the Southern hemisphere to take advantage of the alternating seasons. Each of these options involves discounts that may themselves make it harder to sell clothing at full price in the first place. For this reason, Burberry, a luxury brand, decided to forego the insurance and dispose of their excess rather than cause price seeking across time. Below we will consider explicitly the interaction between the adoption of AI prediction and inventory management.

With AI prediction, an additional choice may present itself. It may be possible to change what is produced based on predicted demand. If the predictions are good enough, there will be little need to worry about too much or too little inventory.[2] To the extreme, better prediction could mean there is no need for insurance practices like warehouse sales at all.

Fire prevention: Another example of insurance is the preparations for fires in buildings. In 2019, as a horrified world watched, Paris' famed Notre Dame cathedral burned. At the time, there was a real fear that the entire structure would be destroyed, wiping away centuries of history and one of France's major cultural icons. But a few hours later, the fire was under control. While the damage was extensive, a remarkable amount of the building was saved including the two front towers. What was lost was the largely wooden roof and a spire that stood atop it. Most critically, and some would say miraculously, no lives were lost.

While no one would have expected market insurance could cover what might be lost from a fire at Notre Dame, other ways of reducing fire damage were also limited. There was no sprinkler system installed throughout the roof because it was too complicated in the old building, potentially increasing fire risk because of the wiring requirements. Moreover, the roof was a lattice built from what was now old wood, so nested that it was referred to as 'The Forest.' Thus, the building itself was not protected by fire-resistant materials the way many newer buildings are.

A considered risk-management system was developed instead. The idea was to rapidly respond to any fire, what we label 'insurance' as in the ability to reduce the damage from a harmful event (what economists would call a bad state) when it occurs. This involved regular training exercises by the Paris fire department, the installation of a 'dry pipe' to bring water in if needed, the assumption that the old oak frame would not burn quickly and, finally, an alert system that allowed for the fire response to be triggered rapidly.

[2] The industry has evolved further to deal with these issues, including attempts at prediction. There is, of course, an entire industry to make sure that the designers' choices of what is fashionable is communicated to the public. That can help sure-up demand to some degree. However, more recently, clothing manufactures have changed their production and distribution so that they can, if high demand is observed, supply the market quickly – in weeks rather than months. Companies like Zara, HM and Forever 21 specialize in such fast fashion. This fashion is low cost but also low quality. It is designed, however, to sell rather than become excess stock. That said, being low cost, the costs of that outcome are correspondingly lower as well and so fashion waste has grown. It is like 'whack-a-mole.' Try to reduce waste in one dimension and it appears in another way.

Ultimately, the response 'worked' in terms of saving the entire building and lives but the fact that the fire was so large was because the alert system was flawed. It took 30 minutes from the time a fire alert was seen by the security guard on duty before the fire was located and responders informed.[3] This was a combination of inexperience of the guard on duty and the confusing nature of the alert system. By the time firefighters actually arrived, it was out of control. Not only that, the dry pipe system had a previously undetected leak and so water had to be drawn from the Seine.[4] The ultimate conclusion was that the fire management plan was sound but had been poorly designed and required more redundancies.

How we deal with the possibility of fire is a microcosm of the options available to manage risk. Notre Dame is an example of an insurance strategy, which involves dealing with the consequences of bad outcomes when they occur. They put in resources, not to prevent fire or minimize the probability that a fire occurred, but instead, to make sure they could contain any fire that did arise. Ultimately, that system was itself not perfect but what did work worked well enough to prevent the worst outcomes. By contrast, this option is not used for most buildings. Instead, a sprinkler system might be installed to reduce the damage any fire might cause. But the goal is the same. Fires themselves are not prevented but their consequences are directly dealt with.

2.2 Examples of protection

We can also prevent fire by (a) choosing building materials that cannot catch fire or (b) choosing what activities can be conducted in buildings to minimize the probability of a fire. In the first case, we make investments that reduce the probability that a building can catch fire. In the second case, we avoid 'playing with matches' and other activities that can cause fires in the first place. Thus, rather than an insurance response, in managing risk we can choose to protect against the risk directly or by exercising more caution in our activities. Which of these risk management options we choose depends on a number of trade-offs that can be different depending on context.

In our model to follow, we distinguish between insurance (reducing the cost of bad outcomes) and protection (reducing the probability of bad outcomes). Fire management involves both. But our focus here is on the activities that support each type of risk management strategy. As we will see, prediction is often a substitute for those activities so, in order to understand the broad impacts of AI adoption, it is important to identify those activities explicitly.

Brick houses: Consider, for instance, the protection story encapsulated by the children's

[3] https://www.nytimes.com/interactive/2019/07/16/world/europe/notre-dame.html
[4] https://www.phcppros.com/articles/10231-lessons-in-fire-protection-from-notre-dame-cathedral

story, *The Three Little Pigs*. In that story, three pigs, faced with a threat of a big-lunged wolf, build houses made out of straw, wood and bricks respectively. At the end of the day, only the brick house is effective against the wolf which justifies the expense in terms of time, materials and forgone leisure, that pig incurred. The moral of the story is that you have to pay a price to be safe. The related moral is that you should pay that price but, as economists, we cannot as readily accept that conclusion. Costs must always be weighed against benefits. Still, the idea is that once it is decided that a brick house is worth building, then a rule is created: Build brick houses because the chance a wolf might come by is high enough.

Over time, the rule becomes invisible. Brick becomes part of the building code for city, and even if the wolf population disappears, people will keep building their houses from brick. When protection turns into a rule, it may be difficult for organizations to recognize the benefit of AI in reducing the underlying uncertainty.

More generally, we have building codes that precisely specify various measures that must be taken to protect those inside buildings from uncertainty events. These include the aforementioned fire prevention but also damage from weather, foundation security, and other natural phenomena like earthquakes.

Hedgerows: Consider the long-standing protection employed for farming in England – building hedgerows. A hedgerow is a carefully planted set of robust trees and plants that can serve as a wall between fields. This is extremely useful if your field is full of farm animals and you do not want to employ a person to ensure they do not wander off. It is also useful if you do not want a heavy rainfall to cause soil to erode too quickly or if you want to protect crops from strong winds. Given all this protection against risky events, it is not surprising that this practice was the origin of the term, hedging, which evolved to have a broader (insurance) meaning.

But hedgerows come at a cost. By dividing farmland, they make it impossible to use certain farming techniques – including mechanization – that are only efficient for large swathes of land. After World War II, the British government actually gave subsidies for the removal of hedgerows for this purpose although in some cases, that removal was excessive given their role in risk management. Today, there are moves to restore hedgerows led most prominently by the Prince of Wales.[5]

[5]From https://www.washingtonpost.com/graphics/2019/world/british-hedgerows/

> The BBC recently aired a documentary called "Prince, Son and Heir: Charles at 70."
>
> In the show, Prince Charles's two sons discuss their father's passion.
>
> "He loves his hedgelaying," Prince William says.
>
> "Whichever policeman is on duty at the time puts the sledgehammer and ax in the boot of the car," Prince Harry says. "Off they go. They spend two hours wrestling with bushes to try to lay a hedge because he hates fences."

We see many situations where costly investments are made to cover or shelter a decision-maker from risk. Miles of highways are cocooned with guard rails to prevent cars from going down embankments, hills, or into on-coming traffic. Most of these are, fortunately, never used but each allowed a road to be built in a way that might otherwise have not been sufficiently safe given the fallibility of human drivers.

Aircraft life vests: What these protection measures have in common is that they typically generate what look like over-engineered solutions. They are designed to be rated up to a certain set of events – the once-in-a-lifetime storm or the once in a century flood. When those events occur, the engineering looks worthwhile. But, in their absence, there is cause to wonder. For many years, as Levitt & Dubner (2005) pointed out, life vests and rafts on aircrafts – not to mention the safety demonstrations of each – appeared to be wasteful given no aircraft had successfully landed on water. Then, in 2009, Captain Sullenberger landed a USAir plane with no working engines on the Hudson River. Does that one example of a low probability event make the precautionary life vests worth it? It is hard to know. But it could certainly not be concluded that the absence of a possible outcome would cause us to assess the probability of that outcome at zero.

Levitt and Dubner's main point, however, is that while it is often possible when protection activities are employed to measure the likelihood or change in likelihood of underlying uncertainty over time, it is not possible to measure whether the investments made to reduce the probability of a consequence are excessive as the very risk management strategy employed takes away that information. It is entirely possible that too much is being wasted on something that, for other reasons, is no longer a high risk at all.

2.3 Summary

The model below will consider the two, distinct risk-management strategies of insurance and protection. Insurance involves mitigating the downside risk outcomes (say by having a payout in the case of a fire) while protection involves reducing the probability itself of a bad outcome (say by using different building materials).[6] There is no necessary bright line between each and sometimes risk management involves the use of both even though they are ultimately serving the same end. However, what we want to emphasize is that each strategy has potentially different implications as to how we observe and learn about the underlying uncertainty that requires risk management in the first place.

Harry says, "Some come back covered in blood because at some point something he has been cutting has flung up."

[6]This distinction was first made by Ehrlich & Becker (1972).

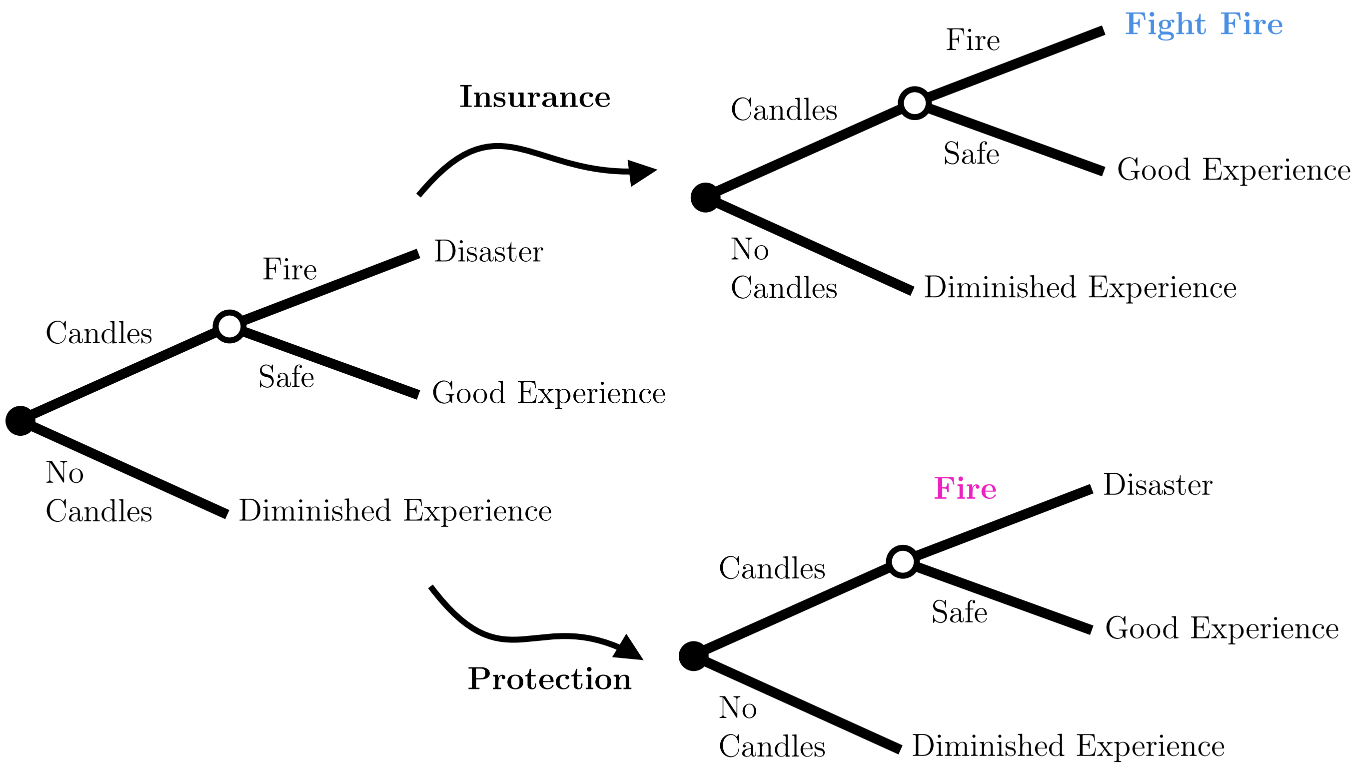

Figure 1: Insurance vs Protection under Fire Risk

The key focus here is how risk-management and then prediction impact decisions. Consider the following example. Fires can happen independently of what people do but, as is readily apparent, fires can also be started by the activity of people. This is why we don't like children or basically anyone else to 'play with matches.' Notre Dame's fire risk would be much lower if people weren't allowed in and it didn't have to accommodate people by having, say, electric or other forms of lighting and temperature control in the building.

When you could visit Notre Dame, worshippers could light a candle. Obviously, candles aren't necessarily the greatest risk for causing a fire in a centuries old cathedral but many buildings prevent activities with open flames. However, for illustration purposes, consider Notre Dame's decision to allow candle lighting or not. The relevant decision tree is depicted in on the left hand side of Figure 1. Allowing candles risks a fire disaster but permits a better experience visiting Notre Dame. Not allowing candles reduces fire risk – actually in the figure it eliminates it – but diminishes the experience. Notre Dame previously allowed candles so must have judged that the additional fire risk was low enough relative to the benefits of having a good experience over a diminished one.

However, as already noted Notre Dame engaged in an insurance risk strategy that involved being ready to fight a fire should it arise and, thereby, mitigate the extent of the disaster that might otherwise occur. The upper right hand side of Figure 1 depicts how this insurance strategy changes the decision tree facing Notre Dame. With insurance in place, the potential downside outcome if a fire arises is reduced and will make it safer to permit candles.

The lower right hand side also shows an alternative protection strategy. That strategy changes the decision tree and makes it more attractive to permit candles. It does so by directly reducing the probability that a fire will occur. With that protection in place, candles are less of a risk. Recall that such protection strategies would involve first resistant building materials. [7]

When insurance and protection are used, decision makers find it less inefficient to always choose the same action. They rely on a rule, and do not respond to the underlying state. When insurance and protection are first implemented, decision-makers might appreciate the risks they are managing and the reasons for the rules. However, as time goes on, the rationale may itself recede. Some of this is by design, as one of the advantages of good risk management is not having to worry about uncertainty. But it is also possible that, as the underlying causes of uncertainty shift, or the trade-offs between alternative options change, the risk management strategy itself becomes sub-optimal. In other words, by virtue of past risk management choices, the uncertainty decision-makers might face is hidden. From the perspective of AI adoption, this means that the prediction problems that might be solved may simply not be defined or even noticed. This is both a detriment to businesses tied to old ways of risk management and an opportunity for others looking to apply AI in new areas. If the uncertainty were known, the potential for AI to provide predictions and enable decisions would be clear.

In the case of fire, AI prediction of fire risk could enable new decisions, beyond the current options of more robust materials, less use of fire, and fast response when a fire occurs. Instead, AI prediction could enable decisions on more types of materials or new types of activities that seemed risky under rules, but aren't when predictions are available. Under rules, it may not be clear that the uncertainty limits the scope of decision-making.

3 Basic Model

Our model focuses on decisions under uncertainty; initially, for a single task. An agent has a choice between two actions, $a \in \{1, 2\}$. There are two possible states of the world $\theta \in \{1, 2\}$ that determine the payoffs from those actions. If the action choice matches the state, there is a payoff of R while if the action does not match the state, the payoff is $r < R$. The prior

[7] Interestingly, a sprinkler system has an impact on both the risk of a disastrous fire (by enabling an automatic response as soon as fire is detected) and by being an automated clean-up (via the use of water). Thus, some risk management strategies have elements of protection and insurance.

probability that $\theta = 1$ is $p > \frac{1}{2}$. Thus, absent any other information, it is optimal for the agent to choose $a = 1$ always earning an expected payoff of $pR + (1-p)r$. Choosing one action all of the time is what we will term, *following a rule*.

By contrast, if the agent has perfect information regarding the state, then they will match the action with the state. This would result in a payoff of R and would be implemented by a state-contingent decision. When the action chosen varies with the predicted state, as it would in the case of perfect information, we term this, *making a decision*.[8] When an agent is following a rule they have no need for a prediction of the state whereas the opposite is true if they are making a decision.

Our focus here is what drives an agent to make a decision rather than follow a rule; in particular, the role better prediction plays in this. This is done in the context of three options the agent can take for mitigating the negative impacts involved in not matching the action with the correct state. The first two are from Ehrlich & Becker (1972) while the contribution of this paper is to add the third so as to analyze the impact of artificial intelligence following Agrawal et al. (2019) but incorporating explicitly risk management strategies used under rules or decisions.

- *Protection*: this is a costly, preemptive action designed to shift the probabilities associated with the more frequent state, $\theta = 1$. For instance, this might be building with strong materials like brick. We assume that at a cost, $C(x)$, which is non-decreasing and convex, the agent can choose a higher level of $p+x$. We assume that $C'(1-p) = \infty$ so that setting $x = 1-p$ is never optimal.

- *Insurance*: this is a costly, preemptive action designed to increase the downside payoff (r) from the risky action. We assume that at a cost, $c(\Delta)$, the downside risky payoff becomes $r + \Delta$. c is convex and non-decreasing. We also assume that $c'(R-r) = \infty$ so that setting $\Delta = R - r$ is never optimal. An example of this would be slack in system or inventories; that is, prepared resources to be able to react in a crisis.

- *Prediction*: this is the acquisition of information that provides a signal of the state of the world so that mismatches can be avoided if the (posterior) probability that the unmatched state arises is too high. Thus, we assume that there is a signal that with probability e accurately reveals the correct state and with probability $1-e$ reveals the incorrect state. If a signal of state 1 is received, then the posterior probability that it is actually state 1 is $p(1) = \frac{ep}{ep+(1-e)(1-p)}$. If a signal of state 2 is received, then the

[8]There is a sense in which choosing an action when you have a perfect prediction is hardly 'deciding.' However, the model below will accommodate prediction errors although assumptions will be made such that it is still an optimal decision to the follow the prediction. Doing this, however, will involve an explicit cost.

posterior probability that it is actually state 2 is $1 - p(2) = \frac{e(1-p)}{(1-e)p+e(1-p)}$. We assume that $e > \frac{1}{2}$ so the signals are informative. In this case, the agent will only choose to follow the signal if the error rate is sufficiently small.

The activities of protection and insurance can be varied. For instance, as noted above, if you were worried about the risk of damage from fire in a building, you can ensure the building is made with nonflammable materials that would shift the probability of a fire spreading (protection) or build in an alarm system that quickly alerted the fire department (insurance). By contrast, you could use an AI to predict whether a fire was likely. Such a prediction would allow you to monitor whether conditions are such that a fire may occur and take preemptive action to stop it from occurring. By predicting when a fire would occur, and enabling a decision of whether to quickly act to put out a fire, this would reduce the need for protection and for insurance.

3.1 Information processing costs

If $(x, \Delta) = (0, 0)$, following the prediction by making a decision generates an expected payoff of:[9]

$$eR + (1-e)r$$

By contrast, not following the prediction (i.e., a rule) yields $pR + (1-p)r$. All other things being equal, following the prediction will be optimal if $e \geq p$.

However, other things are not equal. While a rule involves not making an on-going choice (although it is effectively made when the rule is formed), a decision requires continual attention. In particular, if information in the form of AI prediction is being fed into the decision, then the act of choice involves information processing costs that may be primarily cognitive.

Thus, we will assume that processing information is not costless for the decision-maker. This is captured simply by assuming that if the agent wants to follow the prediction, it costs them λ to do so. In this case, they will choose to make a decision rather than follow a rule if:

$$eR + (1-e)r - \lambda \geq pR + (1-p)r \implies e - p \geq \frac{\lambda}{R-r}.$$

Note that when x and Δ are positive, this inequality becomes $e - (p + x) \geq \frac{\lambda}{R-(r+\Delta)}$. Thus, it is easy to see that as e increases, the returns to making a decision over following a rule rise. By contrast, higher x or Δ reduce the returns to decision-making.

[9] The expected payoff is $(ep+(1-e)(1-p))(p(1)R+(1-p(1))r)+(e(1-p)+(1-e)p)((1-p(2))R+p(2)r)$ which collapses to $epR + (1-e)(1-p)r + e(1-p)R + (1-e)pr = eR + (1-e)r$.

3.2 No Prediction

Without prediction, because p is assumed to be greater than $\frac{1}{2}$, it is optimal for the agent to follow the rule of setting $a = 1$. However, they can engage in preemptive expenditures in self-insurance and self-protection. Thus, the agent solves:

$$\max_{x,\Delta}(p + x)R + (1 - p - x)(r + \Delta) - C(x) - c(\Delta)$$

Note that the mixed partial derivative of this objective function with respect to (x, Δ) is -1 so self-insurance and self-protection are substitutes. Given our assumptions on the costs of these actions, there is an interior solution which we denote by $(x^*_{RULE}, \Delta^*_{RULE})$. These are characterized by $(R - r - \Delta) = C'(x^*_{RULE})$ and $1 - p - x^*_{RULE} = c'(\Delta^*_{RULE})$, respectively.

3.3 Prediction

Suppose a prediction is available and suppose that the agent follows the prediction and makes a decision. Note that in the absence of insurance and protection (i.e., $\Delta = x = 0$), prediction yields an expected payoff of $eR + (1 - e)r - \lambda$. Compared to a rule without insurance or protection, with an expected payoff of $pR + (1 - p)r$, prediction will be adopted if $(e - p)(R - r) \geq \lambda$ (and a decision will substitute for a rule). Thus, in this baseline case, it is clear that prediction quality, e, must be at least p for the adoption of prediction to be worthwhile.

When insurance and protection are available, then the agent solves:

$$\max_{x,\Delta} eR + (1 - e)(r + \Delta) - C(x) - c(\Delta) - \lambda$$

The following proposition characterizes how this changes the levels of insurance and protection.

Proposition 1 *If the agent moves from a rule to a decision:*

1. *Protection is reduced; i.e.,* $x^*_{RULE} > x^*_{DEC} = 0$.

2. *Insurance is reduced (increased) if* e *is sufficiently high (low); i.e.,* $\Delta^*_{DEC} \leq (>)\Delta^*_{RULE}$ *if* $e \geq (<)p + x^*_{RULE}$

Proof. If the agent makes a decision, $x^*_{DEC} = 0$; i.e., there is no return to protection. For insurance, recall that Δ^*_{RULE} satisfies $p + x^*_{RULE} = 1 - c'(\Delta^*_{RULE})$ while Δ^*_{DEC} satisfies $e = 1 - c'(\Delta^*_{DEC})$ so that $\Delta^*_{DEC} \leq (>)\Delta^*_{RULE}$ if $e \geq (<)p + x^*_{RULE}$. Note also, by the envelope theorem, $t(r+\Delta^*(t)) - c(\Delta^*(t))$ is increasing in t so that $(1-e)(r+\Delta^*_{DEC}) - c(\Delta^*_{DEC}) \leq (>)(1-p-x^*_{RULE})(r+\Delta^*_{RULE}) - c(\Delta^*_{RULE})$ if $e \geq (<) p + x^*_{RULE}$. Thus, for sufficiently high quality prediction, insurance will fall. ■

Thus, for sufficiently high quality prediction, *both* insurance and protection activities are reduced as AI is adopted. Of course, AI will only be adopted if e is sufficiently high. This fact allows us to demonstrate the following:

Corollary 1 *When (x, Δ) are endogenous, a necessary condition for making a decision to be preferred to a rule is that $e > p$.*

Proof. The agent will choose decisions over rules if:

$$(e-p-x^*_{RULE})(R-r) + (1-p-x^*_{RULE})\Delta^*_{RULE} + (1-e)\Delta^*_{DEC}$$

$$\geq \lambda + c(\Delta^*_{DEC}) - c(\Delta^*_{RULE}) - C(x^*_{RULE})$$

Note that at $e = p + x^*_{RULE}$, $\Delta^*_{DEC} = \Delta^*_{RULE}$ and, thus, the incremental payoff from making a decision is $C(x^*_{RULE}) - \lambda$. If this is positive, then making a decision is optimal if $e \geq p + x^*_{RULE}$ which, in turn, implies that insurance will decline following AI adoption. By contrast, when $e = p$, then $\Delta^*_{DEC} \geq \Delta^*_{RULE}$. The payoff from making a decision is:

$$pR + (1-p)(r+\Delta^*_{DEC}) - c(\Delta^*_{DEC}) - \lambda$$

and from following a rule is:

$$(p + x^*_{RULE})R + (1-p-x^*_{RULE})(r+\Delta^*_{RULE}) - c(\Delta^*_{RULE}) - C(x^*_{RULE})$$

By the envelope theorem, following a rule results in a higher payoff. Thus, a necessary condition for making a decision to be optimal is that $e > p$ (as it would also be in the absence of insurance and protection). ■

Therefore, given that insurance will only increase if a decision is made instead of a rule if $e - p \leq x^*_{RULE}$, when protection is relied upon under a rule, this raises the likelihood that insurance may increase following AI adoption rather than decrease. Good predictions reduce the need for protection. With information about the state, it becomes possible to respond without needing to invest to change the likelihood of one state or another. For example, good predictions on where soil might erode reduce the need to build hedges everywhere. With excellent predictions, insurance is less useful. Firms can make inventory decisions confident that they will be correct. However, if predictions are good enough to enable decision-making instead of rules (but far from perfect), then insurance can be even more valuable. It enables the organization to reduce the risk from a wrong decision.

3.4 AI Adoption

Thusfar, we have considered the choice between rules and decisions assuming the agent was starting from scratch. Because AI prediction is a recent technological development, when considering its adoption, this is undertaken in a context where an agent has already previously chosen self-insurance and self-protection levels and these are fixed. Suppose those levels are x_{RULE} and Δ_{RULE} respectively. Thus, the existing payoff (ignoring sunk insurance and protection costs) is $(p + x_{RULE})R + (1 - p - x_{RULE})(r + \Delta_{RULE})$.

What level of $\{e, \lambda\}$ will cause a switch from a rule to AI adoption and a decision being made?

Proposition 2 *AI prediction will be adopted if* $(e - p - x_{RULE})(R - r - \Delta_{RULE}) \geq \lambda$.

Proof. Given that x and Δ are fixed, this leads to an expected payoff, if AI prediction is used, of $eR + (1 - e)(r + \Delta_{RULE}) - \lambda$. Thus, AI prediction will be adopted if:

$$eR + (1 - e)(r + \Delta_{RULE}) - \lambda \geq (p + x_{RULE})R + (1 - p - x_{RULE})(r + \Delta_{RULE})$$

which gives the condition of the proposition. ■

Note that a necessary condition for AI adoption here is that $e > p + x_{RULE}$. Thus, Proposition 1 shows that AI adoption will result in a reduction in Δ. Moreover, the higher are x_{RULE} and Δ_{RULE}, the lower is the return to adopting AI. In other words, if the actions that can be made to mitigate the waste associated with rules are more effective, the lower are the returns to adoption AI prediction.

3.5 Asymmetric Stakes

The model here presents AI adoption in the context of how prediction might be used to alter how an agent manages risk when facing choices made under uncertainty. Bresnahan (2020) argued that one of the constraints on businesses adopting AI prediction is that they face what he terms 'high stakes loss functions.' In explaining this, he contrasted Amazon's recommendations to customers of what they might purchase with Facebook's predictions of whether content posted is unsafe or offensive. In Amazon's case, Bresnahan argued that the loss function involved low stakes because an incorrect recommendation may lead to a loss sale but not necessarily a lost customer. By contrast, a "false positive" prediction from Facebook's content AI, may lead to inappropriate content being posted that may drive customer churn or, at the very least, controversy. He argued that this is why Amazon relies on its AI for recommendations while Facebook employs a large labor force for content moderation that works alongside its AI.[10]

To explore this within the context of the model here we need to enrich the payoff space. The model here treats both states as symmetric in terms of their impact on payoffs; that is, match the right state to the right action and the agent earns R while a mismatch earns r or more broadly, $r + \Delta$. While for the analysis of a rule this assumption is innocuous, when a decision is being made, it has the effect of creating a symmetry in the stakes – that is, the loss from an error, being $R - (r + \Delta)$ in each case.

Suppose, instead, we introduce asymmetric stakes. We do this in two ways. First, we assume that if $a = 1$ when $\theta = 2$, the agent's uninsured payoff is r_2 while if $a = 2$ when $\theta = 1$, it is r_1. Second, we assume that insurance expenditures are state directed. That is, the agent chooses $\{\Delta_1, \Delta_2\}$ at cost $c(\Delta_1) + c(\Delta_2)$ where Δ_θ are insurance levels if there is a mismatched action for θ (i.e., $a \neq \theta$). To focus on the stakes issue we assume here that $p = \frac{1}{2}$. c is assumed to be convex. For the choice of protection, we will now assume that $C(.)$ is a function of $|x|$ to reflect the notion that x can be negative and shift probabilities away from state 1.

With this set-up we can show the following:

Proposition 3 *Suppose that $r_1 < r_2$. When choosing a rule, it is optimal to set $a = 1$, $x^*_{RULE} > 0$ and $\Delta^*_{2,RULE} > \Delta^*_{1,RULE} = 0$. When choosing to make a decision, the agent chooses $x^*_{DEC} \leq 0$ and $\Delta^*_{2,DEC} \geq \Delta^*_{1,DEC} > 0$. An increase in the stakes, $R - r_1$, leads to a decrease in x^*_{DEC} and $\Delta^*_{1,DEC}$, an increase in $\Delta^*_{2,DEC}$ and a reduction in the returns to AI adoption.*

Proof. Note that $r_1 < r_2$ implies that choosing $a = 2$ involves generically higher stakes than $a = 1$; that is, a higher potential loss $R - r_1$ from a mismatch. Given these changes, if the agent adopts a rule with $a = 1$ always their expected payoff is $(\frac{1}{2} + x)R + (\frac{1}{2} - x)(r_2 + \Delta_2) - c(\Delta_1) - c(\Delta_2) - C(|x|)$. Clearly, with this rule it is optimal to set $\Delta_1 = 0$.

Now consider the impact of AI prediction on decision-making. In this case, the agent's expected payoff is:

$$(e(\tfrac{1}{2} + x) + (1 - e)(\tfrac{1}{2} - x))(\tfrac{e(\frac{1}{2}+x)}{e(\frac{1}{2}+x)+(1-e)(\frac{1}{2}-x)}R + (1 - \tfrac{e(\frac{1}{2}+x)}{e(\frac{1}{2}+x)+(1-e)(\frac{1}{2}-x)})(r_2 + \Delta_2)$$

[10]See also Athey et al. (2020) for a discussion of when AI's versus people have priority in making decisions.

$$+(e(\tfrac{1}{2}-x)+(1-e)(\tfrac{1}{2}+x))(\frac{e(\frac{1}{2}-x)}{e(\frac{1}{2}-x)+(1-e)(\frac{1}{2}+x)}R+(1-\frac{e(\frac{1}{2}-x)}{e(\frac{1}{2}-x)+(1-e)(\frac{1}{2}+x)})(r_1+\Delta_1)$$

$$-c(\Delta_1)-c(\Delta_2)-C(|x|)-\lambda$$

$$=eR+(1-e)((\tfrac{1}{2}-x)(r_2+\Delta_2)+(\tfrac{1}{2}+x)(r_1+\Delta_1))-c(\Delta_1)-c(\Delta_2)-C(|x|)-\lambda$$

Maximizing this with respect to $\{x, \Delta_1, \Delta_2\}$ gives $x^*_{DEC} \leq 0$ and $\Delta^*_2 \geq \Delta^*_1 > 0$.

Note that the payoff under AI adoption is supermodular in $\{r_1, x, \Delta_1, -\Delta_2\}$. Thus, as r_1 falls (i.e., stakes rise), x^*_{DEC} falls, as does Δ^*_1 while Δ^*_2 rises. ■

When following a rule, the agent is at risk of realizing a payoff of r_2 when they choose $a = 1$ always. Thus, they invest in insurance (Δ_2) and protection (x) to manage that risk. By contrast, when making a decision, they are at risk of either r_1 or r_2. Thus, there is positive demand for insurance on either outcome. However, as $r_1 < r_2$, there is an incentive to shift the probabilities of each state towards $\theta = 2$ being more likely. Note also that, for insurance on a $\theta = 2$ outcome, Δ^*_2 will be higher with a decision than a rule case if and only if $(1-e)(\frac{1}{2} - x^*_{DEC}) > \frac{1}{2} - x^*_{RULE}$ or $e < \frac{x^*_{RULE} - x^*_{DEC}}{\frac{1}{2} - x^*_{DEC}}$. Interestingly, for insurance on $\theta = 2$, this is now a complement with greater protection in the form of a more negative x.

The example of fire protection also shows the impact of asymmetric stakes. The consequences of a fire can be devastating. This increases the benefit of investing in protection. An AI might reduce the chance of a catastrophic fire by predicting it before it occurs, but there remains a real risk of a bad outcome. Instead, a rule of using fire-preventing materials might be better.

Finally, note that Proposition 3 confirms our intuition that when stakes are higher, there is a reduced incentive to move away from a rule to a decision using prediction. This confirms, for instance, the contrast between Amazon and Facebook discussed earlier. Note, however, that Facebook does actually use AI prediction to assist in identifying unsafe content. Proposition 3 shows that because of this, they will want to reduce the probability of that occurring – i.e., engage in more protection against $\theta = 1$ when choosing $a = 2$ (akin to allowing content to be posted). Thus, one view of Facebook's employment of thousands of content moderators is that (a) they choose to rely on AI prediction and (b) they decided not to have a rule that led to very little content being posted.

3.6 Empirical implications

What this means is that AI adoption needs to be understood in the context of existing methods of mitigating the consequences of uncertainty. The first aspect of this follows directly, and perhaps obviously, from the viewpoint of AI as prediction: AI will be adopted in decisions where there is a clear benefit from reduced uncertainty. The other aspects require an understanding of rules, protection, and insurance. We expect AI to reduce the need for rules and increase the ability to use decisions. Decision-making will become a more important aspect of work, as Deming (2021) showed is already happening. With AI's diffusion, we expect this to accelerate. Furthermore, we expect AI to reduce the need for protection-related investments.[11] There will be less need to address uncertainty in advance, and more need to respond as things are happening. For insurance, the implications are ambiguous and depend on the quality of the predictions and whether the AI enables decisions over rules.

4 Multiple Tasks

Up until this point, we have considered a model of a single task and the impact of AI on whether that task is guided according to rules (that are state invariant) or decisions (that are not). The literature that evaluates the labor market impacts of AI typically envisages that while AI adoption is at the task level, there are a number of tasks that combine to produce the output of any firm. Those tasks are usually assumed to be complements. Thus, to the extent that the adoption of AI in one task raises productivity of that task, the marginal product of other tasks are similarly enhanced. In this respect, substitution of capital for labor in one task can lead to an improvement in wages and the demand for labor in other tasks.

The model here similarly involves the outcome that AI adoption, when it occurs, is productivity enhancing in that the expected payoff from the task is improved. However, what happens if there is more than one task and those tasks are interrelated in that there are benefits to aligning the decisions across tasks (see Agrawal et al. (2021)). For instance, as well as there being benefits to matching an action to a state, there may be benefits in matching (or aligning) actions across tasks. One way to achieve such alignment is via explicit coordination (say, through communication) but this involves its own costs. Moreover, as we will show, those costs vary depending on whether AI is adopted or not. Indeed, despite there being complements amongst tasks, we highlight the possibility that AI adoption reduces rather than enhances the productivity of interrelated tasks.

[11]It is possible that the number of tasks increases substantially as a result of better prediction. Thus, AI would reduce the proportion of tasks that use insurance or protection but still increase the total amount of insurance and protection.

4.1 The AI Bullwhip Effect

Suppose there are two tasks. The first, that we have already analysed, is chosen by agent A who selects a. The second is chosen by agent B who selects $b \in \{1, 2\}$. Both agents choose their actions to maximize total joint payoff so there are no incentive issues.

Should $a = b$, then a benefit of Γ is conferred but should $a \neq b$, the incremental payoff is $\gamma < \Gamma$. We also assume that at a cost of Θ, a can be communicated to B prior to b being chosen. In this situation, Γ is always realized.

First, suppose that A chooses to follow a rule whereby $a = 1$ always. In this case, B will always find it optimal to choose $b = 1$. In this case, the expected payoff will be:

$$(p + x)R + (1 - p - x)(r + \Delta) + \Gamma - c(\Delta) - C(x)$$

Note that because A follows a rule, B can align their action without any communication/coordination. This structure implies that x^*_{RULE} and Δ^*_{RULE} are not changed from the baseline model.

Now consider what happens should A choose to make a decision and follow the AI prediction. What should B do? In the absence of communication, $a = 1$ which arises with probability $ep + (1 - e)(1 - p)$ as compared with $e(1 - p) + (1 - e)p$ for $a = 2$. Note that $ep + (1 - e)(1 - p) > e(1 - p) + (1 - e)p \implies (e - \frac{1}{2})(p - \frac{1}{2}) > 0$ which is always true. Thus, $a = 1$ is A's most likely action and so it is optimal for B to set $b = 1$ always. In this case, the expected payoff is:

$$eR + (1-e)(r+\Delta) + (e(p+x) + (1-e)(1-p-x))\Gamma + (e(1-p-x) + (1-e)(p+x))\gamma - c(\Delta) - C(x) - \lambda$$

Note that while Δ^*_{DEC} is the same in the basic model, here, there is an extra incentive to choose $x^*_{DEC} > 0$ as this increases the probability that Γ rather than γ is realized.

The key feature of AI adoption in a multi-task environment is that the performance of the second interrelated task is degraded. While a rule leads to that task generating Γ always, when AI is adopted, that falls to $\Gamma - (e - (p + x)(2e - 1))(\Gamma - \gamma)$. This illustrates a *bullwhip effect* associated with AI. Because AI adoption means that A varies their action, this makes it more difficult for B to align with that choice which means that the full alignment benefits are not realized. As noted earlier, protection plays a role in mitigating these effects by reducing the probability that A receives a prediction that $\theta = 2$. In the appendix, we provide a canonical example of multi-decision implications where the decisions are the price and quantity decisions of a firm that can use inventories as a form of insurance. Inventories play an insurance role in that they reduce the consequences that arise if there are shortages as a result of difficulties to predict demand. Better prediction may then reduce the need for inventories but, as we show, this also depends on how tightly coordinated supply decisions are with information regarding demand. These issues are further developed in Gans (2022).

4.2 Managing alignment costs

There are two broad ways in which alignment costs that come from adopting AI can be managed. The first is to invest in communication. If this is done, then A and B are always aligned and Γ is realized. Thus, the cost of adopting AI in a multitask setting is limited to Θ; the cost of communication.

The second way is to invest in insurance for B's task. Thus, suppose that, at a cost of $C_B(\Delta_B)$, if there is misalignment the incremental payoff is increased by Δ_B to $\gamma + \Delta_B$. The choice of Δ_B will be determined by the first order condition: $e - (p + x)(2e - 1) = C'_B(\Delta^*_B)$. Notice that Δ_B and x are complements. This is another manifestation of the AI bullwhip effect that in a multitask environment, adopting AI leads to more investments in insurance and protection than in a single task environment.

The investments in insurance (Δ_B) reduce the loss from misalignment. Note, however, that such investments are not required if communication investments are made. Those communication investments, in this model, eliminate the AI bullwhip effect.

Overall, multiple tasks generate a barrier to AI adoption. Without communication, the need to coordinate creates incentives to continue to follow a rule even if it is clear that the decision might be wrong given the state. If AI is adopted, it can create incentives to invest in insurance beyond what would be needed absent the AI. In this case, the insurance is to mitigate the negative consequences for the second decision-maker whose objective is alignment, rather than just matching the state.

Empirically, this suggests that companies are likely to improve communications processes when adopting AI, particularly in settings where bullwhip effects are anticipated.

5 Conclusion

In this paper, we have presented a model of decision-making uncertainty and examined three distinct risk-mitigation tools. Protection involves reducing the chance of a negative state occurring, just like the hedgerows throughout England that protect against soil erosion. Insurance involves mitigating the consequences of a negative state after it occurs, such as a quick response plan should a fire break out. Prediction involves providing information about the state, and enabling different actions depending on the expected state. Today's AI is best understood as a prediction technology, so we have focused on how better prediction through AI affects the other risk-mitigation tools.

We have shown that better prediction reduces the need for protection. Instead of reducing the chance of one state or another, AI enables the decision-maker to respond optimally given the state. When the AI becomes good enough, it will also reduce the need for insurance. Prediction with AI will generate more decisions and fewer rules. Decision-making will be an increasingly important aspect of work. We have also demonstrated that coordination between decisions within organizations is especially important with prediction, and so within-organization communication is likely to increase.

One direction for future research is to estimate the impact of AI adoption on insurance and protection. For example, a study of industries that includes some companies that adopt AI for decision-making and others that continue with rules might lend itself to a difference-in-differences type of estimation that focuses on the change in investments in insurance and protection before versus after the adoption of AI relative to similar companies that do not adopt or that adopt significantly later. Another direction for future research concerns an empirical analysis of alignment costs. In this case, one could measure the investment in communication at companies before versus after AIs are adopted in Bullwhip-type settings compared to communication investments in similar companies with Bullwhip potential that do not adopt AIs.

Overall, our approach is to focus on AI as it has been developed over the past decade or so. It is prediction technology. As such, in contrast to the existing literature which emphasizes how AI will be able to perform work as well as people and replace labor, we examine how organizations will be able to change the systems they use to address uncertainty.

References

Acemoglu, D. (2021). *Harms of AI*. Technical report, National Bureau of Economic Research.

Acemoglu, D. & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.

Agrawal, A., Gans, J., & Goldfarb, A. (2018a). Prediction, judgment, and complexity: A theory of decision-making and artificial intelligence. In *The Economics of Artificial Intelligence: An Agenda* (pp. 89–110). University of Chicago Press.

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

Athey, S. C., Bryan, K. A., & Gans, J. S. (2020). The allocation of decision authority to human and artificial intelligence. In *AEA Papers and Proceedings*, volume 110 (pp. 80–84).

Bresnahan, T. (2020). *Artificial intelligence technologies and aggregate growth prospects*. Technical report, Stanford University.

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).

Deming, D. (2021). The growing importance of decision-making on the job. *NBER WP 28733*.

Ehrlich, I. & Becker, G. S. (1972). Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4), 623–648.

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).

Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., et al. (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, 116(14), 6531–6539.

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

Levitt, S. D. & Dubner, S. J. (2005). *Freakonomics*. B DE BOOKS.

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).

Rock, D. (2019). Engineering value: The returns to technological talent and investments in artificial intelligence. *Available at SSRN 3427412*.

A Appendix

A.1 A canonical example

To provide an example of how the adoption of AI in one task can impact upon other tasks as well as on risk management choices, we examine a situation of a firm choosing price and quantity under uncertainty. Consider a firm facing uncertain demand. Specifically, for price (P) and quantity (Q), (inverse) demand is $\theta - Q$ where we assume that $\theta \in \{1, 2\}$ as before. The firm has unit production costs of w per unit. If θ was known to all in the firm, $P^* = \frac{1}{2}(\theta + w)$, $Q = \frac{1}{2}(\theta - w)$ and $\pi = (P - w)Q = \frac{1}{4}(\theta - w)^2$.

We assume that the pricing decision is handled by agent A and the quantity decision is handled by agent B. To align with the current model, we constrain the pricing and quantity choices to those that generate the profit maximising outcomes with certainty. That is, $a \in \{\frac{1}{2}(1 + w), \frac{1}{2}(2 + w)\}$ while $b \in \{\frac{1}{2}(1 - w), \frac{1}{2}(2 - w)\}$.[12] Given this, we can calculate the profits that arise from various mismatch scenarios as summarized in Table 1.

Note there that there are two broad errors. If you produce too high a quantity, you have wasteful expenditures. If you price high when demand is low, under the assumptions here you cannot sell any units while you incur wasteful expenditures as a result. By contrast, pricing low when you can price high, means missed sales even when you set $b = Q_2^*$.

Table 1: **Scenarios and Mismatches**

State	a	b	Expected Profit
$\theta = 1$	P_1^*	Q_1^*	$\frac{1}{4}(1-w)^2$
	P_1^*	Q_2^*	$\frac{1}{4}(1-w)^2 - w\frac{1}{2}$
	P_2^*	Q_1^*	$-\frac{1-w}{2}w$
	P_2^*	Q_2^*	$-\frac{2-w}{2}w$
$\theta = 2$	P_2^*	Q_2^*	$\frac{1}{4}(2-w)^2$
	P_2^*	Q_1^*	$\frac{1}{4}(2-w)(1-w)$
	P_1^*	Q_2^*	$\frac{1}{4}(2-w)(1-w)$
	P_1^*	Q_1^*	$\frac{1}{4}(1-w)^2$

In the absence of AI prediction of demand, regardless of the state, there is a single price and quantity chosen. There are four possible pairs here. The following proposition characterizes what the optimal rule is depending upon the exogenous parameters of (p, w).

Proposition 4 *In the absence of AI prediction of demand, it is optimal to choose:*

[12]In reality, these will not be the prices and quantities chosen when there is uncertainty over demand (Lim (1980)). However, the main point here will be illustrated well by this constrained model. The general problem is examined by Gans (2022).

1. $\{P_1^*, Q_1^*\}$ *if* $p \geq \max\{\frac{1-w}{1+w}, \frac{3-2w}{4-w^2}\}$ *with expected profit* $\frac{1}{4}(1-w)^2$*;*

2. $\{P_2^*, Q_2^*\}$ *if* $p < \frac{3-2w}{4-w^2}$ *for* $w < 0.210756$ *and* $p > \frac{2-w}{3-w-w^2}$ *for* $w < 0.210756$ *with expected profit* $\frac{1}{4}(w-2)(p(w+2)+w-2)$*;*

3. $\{P_1^*, Q_2^*\}$ *if* $\frac{2-w}{3-w-w^2} < p < \frac{1-w}{1+w}$ *with expected profit* $\frac{1}{4}\left(-p(w+1)+w^2-3w+2\right)$*.*

Proof. First, note that $\{P_1^*, Q_1^*\}$ is preferred to $\{P_2^*, Q_2^*\}$ for $p \geq p_{11>22} \equiv \frac{3-2w}{4-w^2}$ and is preferred to $\{P_1^*, Q_2^*\}$ for $p \geq p_{11>12} \equiv \frac{1-w}{1+w}$. Second, $\{P_1^*, Q_2^*\}$ is preferred to $\{P_2^*, Q_2^*\}$ for $p > p_{12>22} \equiv \frac{2-w}{3-w-w^2}$. Note that these thresholds (i.e., $p_{11>22} = p_{11>12} = p_{12>22}$) all coincide for $w = 0.210756$. When $w < 0.210756$, $p_{11>12} > p_{11>22} > p_{12>22}$ while for $w > 0.210156$, $p_{12>22} > p_{11>22} > p_{11>12}$. This confirms the rankings in the proposition.

Second, note that profits under $\{P_2^*, Q_1^*\}$ only exceeds profits under $\{P_1^*, Q_2^*\}$ if $1-2w-w^2 < 0$ or $w > \sqrt{2}-1 (> 0.210756)$. Note also that $\{P_1^*, Q_2^*\}$ is dominated by $\{P_1^*, Q_1^*\}$ if $p \geq p_{11>21} \equiv \frac{1}{2+w}$. Note, however, that if $w = \sqrt{2}-1$, $p_{11>21} = 0.414214$. At this point, the payoff (0.232233) from $\{P_2^*, Q_2^*\}$ exceeds that of both $\{P_1^*, Q_1^*\}$ and $\{P_2^*, Q_1^*\}$ (0.0857863) but for higher $p > \frac{1}{2+w}$ the latter payoff is declining while the former is constant. Thus, $\{P_2^*, Q_1^*\}$ is dominated. ∎

Note that $\{P_2^*, Q_1^*\}$ is never chosen as it results in pure waste if $\theta = 1$ and the same outcome as $\{P_1^*, Q_2^*\}$ if $\theta = 2$. The interesting result is that $\{P_1^*, Q_2^*\}$ can be preferred to either 'fit' options where price and quantity are aligned. This is because, while there is excess supply when $\theta = 1$, price is low reducing excess supply while potential excess demand is mitigated when $\theta = 2$ by having a larger quantity available. This means that, in contrast to our baseline model, sometimes a rule with misaligned choices of a and b is preferable to aligned choices.

This outcome only arises for w low. Hence, to focus on cases aligned with the baseline model, we assume here that $w = 0.585786$ which means that $\{P_1^*, Q_1^*\}$ is the optimal rule when p is greater than $\frac{1}{2}$ while $\{P_2^*, Q_2^*\}$ is optimal otherwise; aligning this model as closely as possible with the baseline model.

What happens if A has access to an AI prediction which accurately reveals θ with probability e and is mistaken with probability $1-e$?

1. If A follows the prediction then, if $b = Q_1^*$, expected profit is:

$$e(p\tfrac{1}{4}(1-w)^2 + (1-p)\tfrac{1}{4}(2-w)(1-w)) + (1-e)(p\tfrac{1}{4}(2-w)(1-w) + (1-p)\tfrac{1}{4}(1-w)^2)$$

$$= \tfrac{1}{4}(1-w)(1-e(2p-1)+p-w)$$

while if $b = Q_2^*$, expected profit is:

$$e(p(\tfrac{1}{4}(1-w)^2 - w\tfrac{1}{2}) + (1-p)\tfrac{1}{4}(2-w)^2) + (1-e)(-p\tfrac{2-w}{2}w + (1-p)\tfrac{1}{4}(2-w)(1-w))$$

$$= \tfrac{1}{4}\left((2-w)(1-pw-p-w) - e\left(p\left(w^2-w+1\right)+w-2\right)\right)$$

2. Given this B will choose Q_1^* rather than Q_2^* if $e \geq \frac{1-w-p\left(3-w^2\right)-w+1}{p(w^2+w-1)-1}$.

The right hand side of the above inequality is (weakly) increasing in p and w.

Suppose that $p = 0.6$ and $w = \frac{1}{2}$ so that the optimal rule involves $\{P_1^*, Q_1^*\}$. In this case, if AI is adopted B will choose Q_2^*. Thus, the net payoff to adopting AI is

$$\tfrac{1}{4}\left((2-w)(1-pw-p-w) - e\left(p\left(w^2-w+1\right)+w-2\right)\right) - \tfrac{1}{4}(1-w)^2$$

Thus, AI will be adopted if $e > 0.809524$. Interestingly, when AI is adopted more inventories will be observed than when a rule is followed. This is in contrast to the result of Milgrom & Roberts (1988) that information provision of this kind is a substitute with inventories.

What is going on here is the inventories can play a role when there is uncertainty in providing insurance but also in buffering the AI bullwhip effect. The net effect on inventories is a function of which role is more salient.

A.2 A Remark on Task Production Functions

The impact of AI adoption, therefore, impacts both on the productivity of individual tasks and also on their interaction. In particular, it is possible that AI adoption increases the elasticity of substitution between tasks muting the effects of changes in the marginal product of individual tasks.