NBER WORKING PAPER SERIES

DEMAND FOR ARTIFICIAL INTELLIGENCE IN SETTLEMENT NEGOTIATIONS

Joshua S. Gans

Working Paper 32685
http://www.nber.org/papers/w32685

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024

Joshua Gans has drawn on the findings of his research for both compensated speaking engagements and consulting engagements. He has written the books Prediction Machines, Power & Prediction, and Innovation + Equality on the economics of AI for which he receives royalties. He is also chief economist of the Creative Destruction Lab, a University of Toronto-based program that helps seed stage companies, from which he receives compensation. He conducts consulting on anti-trust and intellectual property matters. He also has equity and advisory relationships with a number of startup firms. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Demand for Artificial Intelligence in Settlement Negotiations
Joshua S. Gans
NBER Working Paper No. 32685
July 2024
JEL No. K41,O31

ABSTRACT

When AI prediction substantially resolves trial uncertainty, a party purchasing AI prediction will disclose it if it is in their favour and not otherwise, signalling the outcome to the other party. Thus, the trial outcome becomes common knowledge. However, this implies that the parties will settle rather than purchase the AI prediction. When parties have differing prior beliefs regarding trial outcomes, these differences are only resolved if the AI prediction is purchased and utilised. In this case, AI will be purchased in equilibrium. Different trial cost allocation rules awarding all costs to the losing party (the English Rule) or having each party bear their own costs (the American Rule) can impact the demand for AI for settlement negotiations, but how this occurs interacts with the expectations regarding whether a settlement will occur or not in AI's absence.

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

1 Introduction

Artificial intelligence (or AI) has been advancing due to substantial improvements in machine learning since 2012. One of the professions that AI may impact is legal work (Susskind and Susskind (2015), Felten et al. (2021)). While such forecasts have focussed on the complexity of summarising and evaluating legal documents, startups are emerging whose purpose is to improve the interpretations and expectations of legal outcomes themselves (Aidid and Alarie, 2023). For instance, Toronto startup, Blue J Legal, has trained AI algorithms on available tax law and precedents, launching products that allow tax accounts, lawyers and even judges to provide assessments of the likely legal outcomes in particular circumstances (Alarie et al. (2016), Alarie et al. (2018)). And there are many competitors moving into this space beyond taxation law (Scholosser, 2024), including incumbents such as Thomson Reuters, Lexis Nexus and Bloomberg Law.

At the core of these products is a simple fact about current AI based on machine learning; that it is an advance in the computational statistics of prediction (Agrawal et al., 2018a). This means that AI products are prediction machines, but done at a much larger corpus of data and capable of returning high-fidelity predictions based, often, on natural language inputs. While the precise nature of evidence may, say, differ from case to case, AI prediction is capable of resolving considerable pre-trial uncertainty about the application of the law and the relevance of precedents so provide litigants a clearer prediction of trial outcomes both on liability and damages. As Niblett and Yoon (2024) argue, there are important disagreements that can potentially be resolved by AI.

The goal of this paper is to model the demand for AI products in the context of litigation settlement negotiations. As those negotiations take place in the shadow of forecasts of trial outcomes, the potential of AI to shed light on those outcomes and create common knowledge amongst parties also implies that such products may be valuable to those parties in order to generate a pre-trial settlement.

In the process of providing this formal analysis, an important subtly emerges. Significantly, if parties can use AI prediction to settle outcomes, might they also have incentives to settle before paying for such AI predictions and eliminate the demand of the AI provider? If AI products are costly, parties have an additional incentive to avoid paying for the cost of those products by settling in advance of incurring those costs. Section 2 shows this possibility has important implications for AI providers and the pricing of AI products in this area. Section 3 then examines the case where parties have differing priors about trial outcomes. Interestingly, even in this case, it is shown that only one of the parties will have an incentive to purchase the AI product in equilibrium. This is because they will disclose the

AI product's prediction when it is favourable while a failure to disclose will send a signal of the prediction to the other party. Finally, section 4 considers the implications of alternative cost allocation rules following a trial – the American rule (used in the baseline model) and the English rule. It is demonstrated that differing rules have different implications for the demand for AI legal products.

2 Baseline Model

Consider the canonical model of settlement negotiations in the face of legal litigation (Shavell, 1982). In that model, a Plaintiff (P) and defendant (D) negotiate over a settlement payment, t. All parties are risk-neutral, and they have common knowledge of their costs and priors. If the matter goes to trial, then with probability p, P wins a judgment of J from D; otherwise, no payment is made. The trial costs P, c_P and D, c_D. It is assumed that $pJ \geq c_P$, so that P will go to trial if there is no settlement.

2.1 Settlement Negotiations without AI

The timeline for the settlement game is as follows:

1. P and D negotiate over a settlement payment (t) from D to P. Whether a settlement occurs and the payment amount is determined by the bilateral Nash bargaining solution.

2. If an agreement is reached, the game ends with P receiving t from D.

3. If no agreement is reached, the trial takes place, and both parties pay their respective trial costs (i.e., costs are allocated according to the American Rule[1]), and uncertainty over the judgment is resolved. If the judgment is in P's favour, they receive J from D. Otherwise, no transfers are made.

Under Nash bargaining, the parties settle because the outcome of a judgment is a mere transfer, and the trial costs are positive. The settlement payment, $\hat{t}$, is given by:

$$\hat{t} - (pJ - c_P) = -\hat{t} - (-pJ - c_D) \Longrightarrow \hat{t} = pJ + \tfrac{1}{2}(c_D - c_P)$$

Observe that the settlement amount is higher as P's prospects increase (p) and as the relative costs of going to trial move in P's favour ($c_D - c_P$).

[1]Below the English Rule where the loser pays all costs is considered.

2.2 Settlement Negotiations with AI

How does AI impact these negotiations? What AI potentially allows parties to do is to predict the outcomes of events they are uncertain about. In the context of the model presented here, both parties are uncertain whether the judgment will be in the plaintiff's or defendant's favour. While, in practice, AI may reduce this uncertainty by some amount, here we investigate a limit case where the AI's prediction is a perfect signal of the trial's outcome. It will become apparent that the main qualitative results would be unchanged if the AI's prediction were imperfect.[2]

The AI prediction is supplied by an independent firm (a monopolist). The purchasing party can disclose that prediction; in which case it becomes common knowledge. In other words, the report from the AI supplier is verifiable and credible.

The availability of a trial prediction means that, prior to the game in the previous subsection, the following timeline takes place:[3]

1. The AI provider posts a price f for the AI prediction.
2. P and D negotiate over whether to settle prior to either purchasing the AI. The negotiation outcomes are determined by the Nash bargaining solution.
3. If an agreement is reached, the game ends with P receiving s from D.
4. If no agreement is reached, P and D independently choose whether to purchase the AI.
5. For any party purchasing the AI, uncertainty over the trial outcome is resolved. That party can choose whether to disclose the AI prediction to the other party or not.
6. The original settlement game proceeds based on the (possibly) new information sets.

To understand the outcomes of this extended game, we work backwards.

What happens if P purchases the prediction? They will disclose the prediction only if it is favourable to them; that is, if it predicts they will prevail at trial. Given this, P will have an incentive to disclose the prediction in this event, and the settlement amount becomes

[2]The price might change, but the biggest issue is that the beliefs of the parties will not become common knowledge, which would make the analysis more complicated. The assumption of a perfect signal is designed to remove these complications although it should be noted that the demand for an AI with an imperfect signal would be lower than the one model here.

[3]Casey and Niblett (2020) provide a simple model of AI use in settlement litigation. However, their analysis does not include the disclosure stage here, and their purpose is not to examine the demand for AI products themselves.

$t_{AI} = J + \frac{1}{2}(c_D - c_P)$. What happens if the prediction is that P will lose at trial? In this case, P will not disclose the prediction. However, suppose that it was common knowledge that P purchased a prediction. In this case, D will infer from P's non-disclosure that P is predicted to lose. Given this, D will refuse to settle. However, given that P knows they will lose, P will not go to trial, and their payoff is 0. In this situation, P's willingness-to-pay for the prediction is, therefore,

$$v_P = p\left(J + \tfrac{1}{2}(c_D - c_P)\right) - \left(pJ + \tfrac{1}{2}(c_D - c_P)\right) = -(1-p)\tfrac{1}{2}(c_D - c_P)$$

In other words, if $c_D > c_P$, P will not purchase the prediction in this situation. This calculation hinges on P's purchase of the prediction being common knowledge. Below, we will show that this outcome is justified if both parties can observe the price at which predictions are sold.

An analogous argument shows that if D purchases the prediction, it will only be disclosed if it predicts that P will lose at trial, but that non-disclosure signals the opposite. Thus, D either discloses the AI report, which causes P not to go to trial, or there is a settlement in the amount $t_{AI} = J + \frac{1}{2}(c_D - c_P)$ as P correctly infers that D is predicted to lose at trial. A settlement is negotiated on that basis. Given this, D's willingness to pay for the prediction is, therefore,

$$v_D = -p\left(J + \tfrac{1}{2}(c_D - c_P)\right) - \left(-pJ - \tfrac{1}{2}(c_D - c_P)\right) = (1-p)\tfrac{1}{2}(c_D - c_P)$$

This is only positive if $c_D \geq c_P$. Thus, only one party – the party facing higher relative trial costs – has a higher (positive) willingness to pay for the AI prediction. If the AI supplier sets a price less than this willingness to pay, one of the parties will purchase. Given that we have assumed that the costs, judgment amount, and prior probability of P winning are all common knowledge, both parties will know if the other has purchased the prediction or not at the posted price by the firm. This justifies our earlier assumption that a party's purchase of an AI report is common knowledge.

2.3 Demand for AI in Negotiations

Having established each party's willingness to pay for AI prediction, the negotiations prior to that purchase opportunity can be examined. This brings us to our first main result, which shows that the parties will have an incentive to settle if only to avoid f, the cost of purchasing the AI prediction.

Proposition 1 *Suppose that the AI provider sets a positive price of f for the prediction.*

Then, neither P nor D purchases the prediction.

Proof. First, note that the AI provider will set $f \geq 0$ or it will earn a loss. Second, if $f > \max\{(1-p)\frac{1}{2}(c_D - c_P), -(1-p)\frac{1}{2}(c_D - c_P)\}$, then neither party purchases the AI as it exceeds each's willingness to pay. This implies that a positive price, t, is not possible if $c_D = c_P$. Third, as only one party has a positive willingness to pay, then any t within these bounds might be purchased by one party, and if this occurs, this purchase will be common knowledge.

Following the posting of the price, f, P and D can still negotiate a settlement. If they reach an agreement, then jointly, this price is a cost that can also be saved. In this case, if $c_D < c_P$, the settlement price is given by:

$$\begin{aligned}
&t - \left(p\left(J + \tfrac{1}{2}(c_D - c_P)\right) - f\right) = -t - \left(-p\left(J + \tfrac{1}{2}(c_D - c_P)\right)\right) \\
&\implies t = p\left(J + \tfrac{1}{2}(c_D - c_P)\right) - \tfrac{1}{2}f
\end{aligned}$$

while, if $c_D > c_P$, it is:

$$\begin{aligned}
&t - \left(J + \tfrac{1}{2}(c_D - c_P)\right) = -t - \left(-p\left(J + \tfrac{1}{2}(c_D - c_P)\right) - f\right) \\
&\implies t = p\left(J + \tfrac{1}{2}(c_D - c_P)\right) + \tfrac{1}{2}f
\end{aligned}$$

Thus, the AI's existence causes the settlement payment to change but does not change anything else or reduce the joint surplus accruing to the parties. ■

If an AI provider exists and sets a positive price, the parties have an incentive to avoid this cost in addition to the trial costs. The incidence of this cost depends on which party faces the higher trial cost. However, the result here is a negative one regarding the existence of AI, suggesting that no provider would enter this particular market.

This suggests that if such an AI provider did enter, it would not charge for an AI on a 'per case' basis but, instead, a subscription. In that situation, it would not be clear which party would be on which side of any particular case – especially if it was a law firm. This type of 'software as a service' is currently how AI-legal tech firms operate, and the pricing difficulties driving firms towards this are similar to those explored in other AI settings (see Agrawal et al. (2018b)).

3 Differing Priors

In the baseline model, a trial never takes place. For the same reason, parties would not purchase an AI to predict an individual trial's outcome. To enrich the model and explore the impact of AI on trial incidence, we now turn to consider a model whereby each party holds different priors to the likelihood that P wins at trial. This extension is part of the literature on settlement (e.g., Shavell (1989)) and has been the subject of recent work in the area (e.g., Yildiz (2011)).

Let p_P and p_D be P and D's priors respectively. We assume that $p_P > p_D$; that is, both parties are relatively optimistic about their prospects. In this case, a settlement will only occur if $(p_P - p_D) J \leq c_D + c_P$ and the settlement amount will be given by:

$$\hat{t} - (p_P J - c_P) = -\hat{t} - (-p_D J - c_D)$$
$$\Longrightarrow \hat{t} = \tfrac{1}{2}(p_P + p_D) J + \tfrac{1}{2}(c_D - c_P)$$

Thus, in addition to the expected cost differential, the settlement amount, $\hat{t}$, is increasing the average prior probability that the plaintiff will be successful at trial. AI prediction, by providing certainty, can remove this factor from the settlement amount, making that component J or 0 as the case might be. Interestingly, the prior probabilities of each party will factor into their expected benefit from having uncertainty resolved. The question to be addressed next is whether this is a source of differential demand between the parties.

3.1 Demand for AI

With a common prior, we demonstrated above that the AI provider would set a price for AI prediction that would induce only one party to purchase the prediction as the other party's willingness to pay would be negative. When there are differing priors, it will be demonstrated that both parties may have a positive willingness to pay. Nonetheless, under certain conditions, the AI provider will have an incentive to exploit their differences so that only one (possibly) purchases that AI prediction.

In the baseline model, a settlement was always guaranteed to occur whether AI was purchased or not. With differing priors, a settlement may not occur. The demand for AI depends on whether a settlement will occur. Each distinct case will be considered in turn.

First, suppose that $(p_P - p_D) J \leq c_D + c_P$, which implies that a settlement will always occur even with differing priors. If the AI prediction is purchased by one party, the disclosure arguments with the baseline model carry over to this case. The settlement outcomes are then the same as those in the baseline model when AI is purchased; that is, the AI eliminates any differences in priors. Thus, the willingness to pay for AI is, for P,

$$v_P = p_P\left(J + \tfrac{1}{2}(c_D - c_P)\right) - \left(\tfrac{1}{2}(p_P + p_D) J + \tfrac{1}{2}(c_D - c_P)\right)$$
$$= \tfrac{1}{2}(p_P - p_D) J - (1 - p_P)\tfrac{1}{2}(c_D - c_P)$$

and for D,

$$v_D = -p_D \left(J + \tfrac{1}{2}\left(c_D - c_P\right)\right) - \left(-\tfrac{1}{2}\left(p_P + p_D\right) J - \tfrac{1}{2}\left(c_D - c_P\right)\right)$$
$$= \tfrac{1}{2}\left(p_P - p_D\right) J + (1 - p_D)\tfrac{1}{2}\left(c_D - c_P\right)$$

Note that at least one willingness to pay and perhaps both can be positive (for instance, if $c_D = c_P$, in which case, $v_P = v_D > 0$). Note that $v_P > v_D$ if $(1 - p_D)\left(c_D - c_P\right) < -(1 - p_P)\left(c_D - c_P\right)$ or $c_P > c_D$. That is, one party's willingness to pay is higher than the other based solely on the relative trial costs, as was the case in the baseline model.

These willingnesses to pay are calculated on the assumption that only one party purchases the AI prediction. However, if both v_P and v_D are positive, the AI provider could choose to set $f \leq \min\{v_P, v_D\}$ with the aim of both parties purchasing the prediction. The problem is that if one party knew the other was purchasing the prediction, their willingness to pay for the prediction would fall to zero. This implies that if there were to exist an equilibrium where both parties purchased the AI prediction, each one's willingness to pay would be lower than v_P or v_D derived thus far. Given this, a simple assumption is made that excludes an equilibrium where both purchase the prediction. It is assumed that $2\min\{v_P, v_D\} < \max\{v_P, v_D\}$. This is a sufficient condition for the AI provider to prefer to sell to the party with the highest willingness to pay than to both based on the lower willingness to pay of the two parties.[4]

Turning now to the case where $\left(p_P - p_D\right) J > c_D + c_P$ and there is no pre-trial settlement, note that if an AI is purchased by one party, it now enables a settlement. Assuming that f is set so that only one party buys the AI, the willingnesses to pay for the AI are given by:

$$v_P = p_P \left(J + \tfrac{1}{2}(c_D - c_P)\right) - \left(p_P J - c_P\right) = c_P + p_P \tfrac{1}{2}\left(c_D - c_P\right)$$

$$v_D = -p_D \left(J + \tfrac{1}{2}\left(c_D - c_P\right)\right) - \left(-p_D J - c_D\right) = c_D - p_D \tfrac{1}{2}\left(c_D - c_P\right)$$

In this case, both parties' willingnesses to pay for the AI are positive and do not depend on the prior of the other party. Note that $v_P > v_D$ if and only if $c_P > c_D$. Note that, as for the first case above, this calculation applies if it is assumed that $2\min\{v_P, v_D\} < \max\{v_P, v_D\}$.

[4]This is a sufficient condition because calculating the willingness to pay associated with both parties purchasing the prediction, even if they were to purchase with some probability as part of a mixed strategy, will be lower than the willingness to pay derived on the assumption that one is the only purchaser of the prediction. The full set of equilibria could be characterised but involves a more complex analysis than is warranted here, where the idea is simply to explore whether an equilibrium exists where any of the parties purchase an AI prediction.

3.2 Equilibrium AI purchases

When there are differing priors, will the parties come to a settlement prior to any one of them paying for the AI prediction? The following result characterises the conditions under which an AI purchase will occur in equilibrium.

Proposition 2 *Suppose that $c_P \neq c_D$.*

- *If $(p_P - p_D)\, J \leq c_D + c_P$ and $\frac{1}{2}(p_P - p_D)J < \max\{(\frac{3}{2} - p_P - \frac{1}{2}p_D)(c_D - c_P), (\frac{3}{2} - p_D - \frac{1}{2}p_P)(c_P - c_D)\}$, the AI price (f) will be at most $(p_P - p_D)\left(J + \frac{1}{2}(c_D - c_P)\right)$.*
- *If $(p_P - p_D)\, J > c_D + c_P$ and $c_P < \min\{(2 - \frac{1}{2}p_P - p_D)(c_P - c_D), (1 - \frac{1}{2}p_D - p_P)(c_D - c_P)\}$, the AI price (f) will be at most $\max\{c_P + p_P\frac{1}{2}(c_D - c_P), c_D - p_D\frac{1}{2}(c_D = c_P)\}$*

In each case, one party will purchase the AI, and a settlement will emerge.

Proof. The condition that $c_P \neq c_D$ ensures that $v_P \neq v_D$. The two cases are the one where a settlement will occur without AI prediction and one where it will not otherwise occur. The other conditions associated with each case are assumptions that $2\min\{v_P, v_D\} < \max\{v_P, v_D\}$ ruling out the AI provider from attempting to sell AI prediction to both parties.

First, suppose that $(p_P - p_D)\, J \leq c_D + c_P$ and the parties expect to settle in the absence of AI. If P might otherwise buy the AI prediction (i.e., if $c_D < c_P$), then a settlement will take place prior to the purchase of an AI prediction if the joint expected surplus following an AI purchase is less than 0 (or the joint surplus if a settlement would otherwise take place). That is,

$$\begin{aligned} &-p_P\left(J + \tfrac{1}{2}(c_D - c_P)\right) + f + p_D\left(J + \tfrac{1}{2}(c_D - c_P)\right) \geq 0 \\ &\Rightarrow f \geq (p_P - p_D)\left(J + \tfrac{1}{2}(c_D - c_P)\right) = v_P + v_D \equiv \overline{f} \end{aligned}$$

Thus, the AI provider cannot set f above this amount as this would trigger a settlement instead of a purchase. This is only a constraint on the AI provider if:

$$v_P \geq \overline{f} \Longrightarrow (p_P - p_D)\, J \leq -(1 - p_D)(c_D - c_P)$$

This may or may not hold. Thus, for $c_D < c_P$, $f \leq \min\{\overline{f}, v_P\}$.

On the other hand, if D might otherwise purchase the AI prediction (i.e., if $c_D > c_P$), then a settlement will take place prior to AI prediction if

$$\begin{aligned} &-p_P\left(J + \tfrac{1}{2}(c_D - c_P)\right) + p_D\left(J + \tfrac{1}{2}(c_D - c_P)\right) + f \geq 0 \\ &\Rightarrow f \geq (p_P - p_D)\left(J + \tfrac{1}{2}(c_D - c_P)\right) = \overline{f} \end{aligned}$$

This is the same condition as where P might buy the AI prediction. This is only a constraint on the AI provider if:

$$v_D \geq \overline{f} \Longrightarrow (p_P - p_D)\, J \leq (1 - p_P)\,(c_D - c_P)$$

This may or may not hold. Note, however, that, because $\overline{f} = v_P + v_D$ if $c_D < c_P$, then $v_D < \overline{f} < v_P$ and if $c_D > c_P$, then $v_P < \overline{f} < v_D$. Thus, $f \leq \min\{\overline{f}, v_P, v_D\}$ and there exists a f such that only P or D will purchase the AI. Thus, so long as f is set strictly greater than v_P or v_D, then it is common knowledge that one party will purchase the AI, and hence, a settlement will be forthcoming.

Second, suppose that $(p_P - p_D)\, J > c_D + c_P$ and the parties expect to go to trial in the absence of AI. Will the parties settle prior to either purchasing the AI? Comparing the expected joint surplus from settling with that of going to trial, we have:

$$\begin{aligned}
&p_P \left(J + \tfrac{1}{2}\,(c_D - c_P)\right) - p_D \left(J + \tfrac{1}{2}\,(c_D - c_P)\right) - f \geq p_P J - c_P - p_D J - c_D \\
&\Rightarrow (p_P - p_D)\tfrac{1}{2}\,(c_D - c_P) + (c_P + c_D) = v_D + v_P \geq f
\end{aligned}$$

This places an upper bound on the AI's price, but it will not bind. As v_P and v_D, are both positive, if $f \in (\min\{v_P, v_D\}, \max\{v_P, v_D\}]$, the AI will be purchased by one of the parties and a settlement will take place thereafter. ■

When there are differing priors, AI prediction is purchased in equilibrium. The reason is that those differing priors create a reason why a settlement might be harder to achieve and when there is an AI prediction purchased, uncertainty and, importantly, any disagreement between the parties on the outcome of a trial are resolved. That resolution can only occur if AI prediction is actually purchased, and so it is in equilibrium.

Interestingly, the mechanism by which the AI provider selects one party only to purchase the AI prediction and creates common knowledge that it is purchased only works if there is a sufficient differential between the litigation costs of P and D. If these are equal, then $v_P = v_D$ and, thus, there is no price, f, whereby just one party is guaranteed to purchase the prediction. When $c_D = c_P$, this creates a collective action problem as each party would like the other to purchase the prediction and incur f. The proposition assumes that the cost differential is high enough to rule out this outcome. If this was not the case, deriving the equilibrium with a collective action problem is more complex and is left for future research. One solution also not explored is whether the parties might collectively purchase the prediction from the AI provider in this situation.

4 Alternative Cost Allocation Rules

Two broad rules govern the allocation of costs following a trial. Thus far, we have assumed that both parties are in a Court that follows the American Rule where parties bear their own costs regardless of the trial outcome. The contrasting rule is the English Rule, whereby the party that loses at trial bears all the costs. Here, we examine how our results above change if the English rather than American Rule is used.

If the matter goes to trial and there remains uncertainty over a judgment, P expects to earn $p_P J - (1-p_P)(c_P + c_D)$ while D expects $-p_D(J + c_P + c_D)$. In this case, a settlement will only occur if $(p_P - p_D)\, J \leq (1 - p_P - p_D)(c_D + c_P)$ and the settlement amount will be given by:

$$\hat{t} - (p_P J - (1-p_P)(c_P + c_D)) = -\hat{t} - (-p_D(J + c_P + c_D))$$
$$\Longrightarrow \hat{t} = \tfrac{1}{2}\left(p_P + p_D\right) J - \tfrac{1}{2}(1 - p_P - p_D)\left(c_P + c_D\right)$$

Note that the English Rule makes it *less* likely a settlement will be reached than under the American Rule.

When AI prediction is utilised, it assumed here that it is common knowledge that at least one of the parties has purchased the prediction. Then $\hat{t}_{AI} = J + \frac{1}{2}(c_P + c_D)$ if the AI predicts that P will win at trial and is 0 otherwise.

Once again, there are two cases. First, suppose that $(p_P - p_D)\, J \leq (1 - p_P - p_D)(c_D + c_P)$ and the parties expect to settle in the absence of AI. Then the relevant willingnesses to pay for AI are:

$$v_P = p_P \hat{t}_{AI} - \hat{t}_{AI} = \tfrac{1}{2}(p_P - p_D)J + (1 - p_D)(c_P + c_D)$$

$$v_D = -p_D \hat{t}_{AI} - (-\hat{t}_{AI}) = \tfrac{1}{2}(p_P - p_D)J - (1 - p_P)(c_P + c_D)$$

Note that, $v_P \geq v_D \Leftrightarrow 1 \geq \frac{1}{2}(p_P + p_D)$. Hence, the average beliefs regarding P's likelihood of success drive which party is more willing to pay for the AI. As before, the parties will settle before purchasing AI if $f > v_P + v_D$. So long as one party's willingness to pay is strictly different from the other, then the AI will be purchased if $f \leq \min\{v_P + v_D, v_P, v_D\}$.

Alternatively, suppose that $(p_P - p_D)\, J > (1 - p_P - p_D)(c_D + c_P)$ and the parties do not expect to settle in the absence of AI. Then the relevant willingnesses to pay for AI are:

$$v_P = p_P \hat{t}_{AI} - (p_P J - (1 - p_P)(c_P + c_D)) = (1 - \tfrac{1}{2} p_P)(c_P + c_D)$$

$$v_D = -p_D\hat{t}_{AI} - (-p_D(J + c_P + c_D)) = \tfrac{1}{2}p_D(c_P + c_D)$$

Note that, $v_P \geq v_D \Leftrightarrow 1 \geq \frac{1}{2}(p_P + p_D)$. Hence, again, the average beliefs regarding P's likelihood of success drive which party has the higher willingness to pay for the AI. In this case, both parties have a positive willingness to pay for AI. As before, the parties will settle before purchasing AI if $f > v_P + v_D$. So long as one party's willingness to pay is strictly different from the other, then the AI will be purchased if $f \in (\min\{v_P, v_D\}, \max\{v_P, v_D\}]$.

Given this, we can now compare the outcomes under the English versus the American Rule.

Proposition 3 *If at least one party expects to settle under either rule, the maximum possible price paid for the AI is higher under the English Rule than the American Rule. If neither party expects to settle under both rules, the maximum possible price paid for the AI is lower under the English Rule than the American Rule.*

Proof. First, suppose that under each rule, the parties are expected to settle without AI. Note that, under the English Rule, $\max\{v_D, v_P\}$ is greater than the same calculation under the American Rule. This can be shown by comparing the willingnesses to pay and noting that, under the English Rule, $v_P \geq v_D \Leftrightarrow 1 \geq \frac{1}{2}(p_P + p_D)$. By contrast, as the pre-AI settlement amounts are the same in both cases, $\overline{f} = v_P + v_D$ are the same under both rules. Thus, the maximum price that can be paid in equilibrium for the AI is higher under the English than the American Rule.

Second, suppose that, under each rule, the parties are expected to go to trial in the absence of AI. In this case, it is easy to demonstrate that *both* parties' willingnesses to pay for the AI are lower under the English Rule than the American Rule. Consequently, the maximum price that can be paid in equilibrium for the AI is lower under the English than under the American rule.

Finally, suppose that, under the English Rule, the parties are expected to go to trial in the absence of AI, whereas they are expected to settle under the American Rule. This happens if $c_D + c_P \geq (p_P - p_D)\, J > (1 - p_P - p_D)(c_D + c_P)$. In this case, P's willingness to pay for the AI under the English Rule exceeds that under the American Rule if

$$\begin{aligned}
&(1 - \tfrac{1}{2}p_P)(c_P + c_D) > \tfrac{1}{2}\,(p_P - p_D)\, J - (1 - p_P)\tfrac{1}{2}\,(c_D - c_P) \\
&\Rightarrow 2(1 - p_P)c_D > (p_P - p_D)\, J - (c_P + c_D)
\end{aligned}$$

The LHS of this inequality is positive while the RHS is negative and so this inequality always holds. For D, the comparison is:

$$\frac{1}{2}p_D(c_P + c_D) > \frac{1}{2}(p_P - p_D)J + (1 - p_D)\frac{1}{2}(c_D - c_P)$$
$$\Rightarrow -2(1 - p_D)c_D > (p_P - p_D)J - (c_P + c_D)$$

Note that, in this case, the RHS of this inequality is greater than $-(p_P - p_D)(c_D + c_P)$ while the LHS is less than this if $\frac{1-p_D}{p_P-p_D} > \frac{c_P+c_D}{2c_D}$ which holds for $c_D > c_P$ (as the LHS is > 1 while the RHS < 1). Thus, where the willingness to pay under the American Rule for D exceeds that of P, the willingness to pay is lower under the American Rule. ■

The value of AI prediction is potentially higher under the English Rule precisely because that rule encourages more settlement amongst the parties, and that the prediction is valuable should the parties actually settle. However, suppose neither party were to settle regardless of the cost allocation rule applied. In that case, the English Rule implies that each party's willingness to pay for the AI prediction is lower than that under the American Rule. Thus, AI prediction will likely obtain a higher price if parties anticipate the American Rule is applied.

5 Conclusion

This paper has shown that AI can create conditions that facilitate settlements in pre-trial negotiations. However, the price that would be paid for AI prediction in these settings is constrained by the possibility of settling without actually purchasing the AI prediction. Broadly, this demonstrates that the role of AI prediction in negotiations of this kind can have subtle effects.

There are many other contexts in which similar issues might arise. For instance, a buyer and a seller of a house negotiating over a sale price could use AI to predict the likely value of the real estate itself. Alternatively, when a union negotiates with a firm over a wage outcome, it might use AI prediction to understand whether the firm can afford to pay for a wage increase. In each case, AI can resolve uncertainty and potentially facilitate a resolution of negotiations, but, at the same time, the price the AI provider might receive may be constrained by the disclosure-inference issue discussed in this paper.

Nonetheless, it is important to note that the situation considered here is stylised. While there was uncertainty, there were no information asymmetries. AI prediction could potentially resolve these, leading to another dimension upon which AI influences negotiation outcomes.

References

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

Agrawal, A., Gans, J. S., and Goldfarb, A. (2018b). Human judgment and ai pricing. In *AEA Papers and Proceedings*, volume 108, pages 58–63.

Aidid, A. and Alarie, B. (2023). *The Legal Singularity*. University of Toronto Press.

Alarie, B., Niblett, A., and Yoon, A. H. (2016). Using machine learning to predict outcomes in tax law. *Can. Bus. LJ*, 58:231.

Alarie, B., Niblett, A., and Yoon, A. H. (2018). How artificial intelligence will affect the practice of law. *University of Toronto Law Journal*, 68(supplement 1):106–124.

Casey, A. J. and Niblett, A. (2020). Will robot judges change litigation and settlement outcomes? *A First Look at the Algorithmic Replication of Prior Cases.[Online] Available at: http://dx. doi. org/10.2139/ssrn*, 3633037.

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

Niblett, A. and Yoon, A. (2024). Ai and the nature of disagreement. *Philosophical Transactions of the Royal Society A*, 382(2270):20230162.

Scholosser, K. (2024). Legal-tech startup uses ai to search relevant case history to better predict case outcomes. *GeekWire*.

Shavell, S. (1982). Suit, settlement, and trial: A theoretical analysis under alternative methods for the allocation of legal costs. *The Journal of Legal Studies*, 11(1):55–81.

Shavell, S. (1989). Sharing of information prior to settlement or litigation. *The RAND Journal of Economics*, pages 183–195.

Susskind, R. E. and Susskind, D. (2015). *The future of the professions: How technology will transform the work of human experts*. Oxford University Press.

Yildiz, M. (2011). Bargaining with optimism. *Annu. Rev. Econ.*, 3(1):451–478.