LIRAES WORKING PAPER n°2022-01

Market for Artificial Intelligence in Health Care and Compensation for Medical Errors

Juin 2022

Bertrand Chopard and Olivier Musy†*

*Université Paris Cité, LIRAES, F-75006 Paris, France. Mail : bertrand.chopard@u-paris.fr
†Université Paris Cité, LIRAES, F-75006 Paris, France. Mail : olivier.musy@u-paris.fr

DROIT, ÉCONOMIE, GESTION
Sociétés et Humanités
Université Paris Cité

Market for Artificial Intelligence in Health Care and Compensation for Medical Errors

Bertrand Chopard*and Olivier Musy†

June 2022

Abstract

We study the market for AI systems that are used to help to diagnose and treat diseases, reducing the risk of medical error. Based on a two-firm vertical product differentiation model, we examine how, in the event of patient harm, the amount of the compensation payment, and the division of this compensation between physicians and AI system producers affects both price competition between firms, and the quality (accuracy) of AI systems. One producer sells products with the best-available accuracy. The second sells a system with strictly lower accuracy at a lower price. Specifically, we show that both producers enjoy a positive market share, so long as some patients are diagnosed by physicians who do not use an AI system. The quality of the system is independent of how any compensation payment to the patient is divided between physicians and producers. However, the magnitude of the compensation payment impacts price competition. Increased malpractice pressure leads to lower vertical differentiation, thus encouraging price competition. We also explore the effect of compensation on firms' profits at equilibrium. We conclude by discussing our results with respect to the evolution of the civil liability regime for AI in healthcare.

JEL Codes: I11, L13, K13, K41.

Keywords: Artificial Intelligence, Diagnostic, Duopoly, Liability, Physician, Compensation.

*Université Paris Cité, LIRAES, F-75006 Paris, France. Mail: bertrand.chopard@u-paris.fr
†Université Paris Cité, LIRAES, F-75006 Paris, France. Mail: olivier.musy@u-paris.fr

1 Introduction

Diagnostic errors may cause harm to patients. Fortunately, artificial intelligence (AI) systems[1] can establish more precise diagnoses, especially in the early stages of disease. AI systems do this by analysing data from health records and diagnostic images (Aggarva *et al.*, 2022). For instance, a machine learning algorithm can distinguish between different types of cancer, based on images of biopsied tissue. These AI systems have two main characteristics. First, they are already complex, and it is possible that they will have the capacity for self-learning in the future, even eventually becoming autonomous. Second, they involve a multitude of actors: notably developers, manufacturers, operators, and physicians. These two elements raise questions about the liability of actors toward patients if a medical error occurs. The challenge lies in the fact that AI systems may be the direct or indirect cause of patient harm[2] as physicians are also expected to exercise control over them[3] (Higgins, 2022). As a consequence, it is relevant to study the division of compensation for harm paid by the producers of AI systems and physicians.

For the European Commission, the type of AI system the physician uses should determine the choice of liability regime - either strict, or fault-based (see the European Parliament resolution of 20 October 2020 with recommendations to the Commission on a civil liability regime for AI (2020/2014(INL)). For high-risk autonomous AI systems, the Commission recommends a strict liability regime. In this case, compensation payments are divided equally among actors, for instance between AI system producers and operators. The reason is that patients might find it impossible to prove that harm was the fault of the producer or the operator, meaning that their corresponding liability claim would fail[4].

For lower-risk and non-autonomous AI systems, the European Commission considers that the product liability directive applies. This directive allows the patient to claim compensation for harm, subject to the condition that the AI system is proven to be defective. In such a case, only the producer should compensate the patient. Here, the producer refers to manufacturers, developers, programmers, and service providers. Therefore, physicians could reduce their risk of liability when using AI systems. The reduced compensation

[1] The European Parliament defines artificial intelligence as the capability of a computer program to perform tasks or reasoning processes that we usually associate with intelligence in a human being.

[2] Acemoglu (2022) explores the potential harm AI may create in different areas of the economy and social life.

[3] The European Commission recalls that we should seek to avoid a situation in which harmed patients cannot claim compensation (see the European Commission legal framework on AI regulation dated 21 April 2021).

[4] There is a growing pool of literature on liability rules that should apply to autonomous vehicles (Shavell, 2020; De Chiara et al., 2021; Guerra et al., 2021).

paid by physicians would be at the expense of AI systems' producers.

At the same time, physicians operate under fault-based tort law. If we assume that the physician who operates the AI system is held liable for a failure to diagnose the patient, then that physician will have to compensate the patient for any harm caused. But the physician will also have recourse against the AI system producer under the condition that the system caused the harm, in accordance with the product liability directive. Here again, compensation for harm has to be divided between the physician and the system producer. But, in such a case, there is absolutely no reason why the compensation should be divided equally between these two actors.

Finally, the AI system producer can file a civil liability claim against the physician. In the case of a defective AI system, the producer could claim that it is reasonable to hold the physician liable because he or she must manage the risk associated with the system. For example, the physician may have taken an action that influenced its operation. In other words, the producer would simply claim that the medical outcome depended on how the system was being used.

These elements raise a number of questions from the economic analysis of law. First, how do clinicians react to the existence of such AI systems? In other words, how should we represent demand for them? Is there room for strategic concerns, such as reducing the physician's expected liability, as noted above? How are the quality and price of AI systems affected by: (i) the amount of compensation paid to patients; and (ii) the division of compensation payments between physicians and systems' producers? In order to study the effect of liability on the market for AI systems in medical diagnosis, we build on the standard model of vertical differentiation with perfect information (Gaynor et al., 2015). We consider a two-stage game, where firms simultaneously choose the quality of the AI system (i.e., its diagnostic accuracy), and then compete on price. On the demand side, we consider that physicians can either make a diagnosis based on their own capabilities, or they can pay to obtain a diagnosis by an AI system. This scenario has two effects for physicians. First, it can help to establish a better diagnosis, thereby reducing the risk of medical error. The magnitude of the improvement will depend on the quality of the AI system. Second, it can reduce the degree to which physicians compensate their patients for harm, as this compensation may now be shared between physicians and AI systems' producers.

A broad empirical literature explores the association between malpractice liability and healthcare quality in hospitals (Kim, 2007; Iizuka, 2013; Frakes and Jena, 2016; Bertoli and Grembi, 2018). But, to the best of our knowledge, there is no research that simultaneously examines liability, healthcare quality, and the use of

AI systems. This is because AI systems is a new area of research, and we lack data that would help to study their association with liability in the area of healthcare. It is reasonable to assume that physicians will turn to them in order: (i) to improve the quality of care (thus reducing the likelihood of a medical error); and (ii) to transfer some liability to producers if a medical error occurs. Knowing that there is a risk of liability, physicians would avoid diagnosing some patients themselves, in order to reduce their exposure to lawsuits. In our model, this behaviour, known as defensive medicine[5], also improves the quality of healthcare, as the quality of AI systems is higher than that of physicians. The empirical legal studies cited above suggest that the liability risk, that is, the extent to which physicians face the threat of being sued or forced to pay high levels of compensation, is not significantly associated with healthcare quality in the hospital[6]. The medical liability system would not serve to enhance quality in the hospital, due to the so-called 'deterrence' effect. Consequently, in our analysis, we look at how AI system producers can improve healthcare quality, as firms are likely to be more sensitive to liability costs than physicians. Producers focus on improving the quality of care, due to the threat of being held liable for a defective product, while physicians seek to enhance the quality of diagnosis and treatment through the use of AI systems.

The literature on health liability law provides evidence that some medical malpractice cases are settled pre-trial, often with less financial compensation paid by the hospital to the patient (Danzon et al., 1982, Danzon, 1985a,b, Farber and White, 1991). Furthermore, we know that legal costs differ between out-of-court settlements and trials, as does the amount of financial compensation paid by the hospital to the patient (Avraham, 2007). In the present paper, we do not focus on out-of-court settlements. Rather, we study the division of compensation between actors, and its impact on the market for AI systems, based on different price and quality levels.

Finally, we can also link our study with theoretical analyses of product liability (see Daughety and Reinganum, 2013 for a survey). Product liability can be applied here as physicians and AI system producers interact as buyers and sellers. In line with the literature, we consider both market and legal incentives available to agents. Typically, in our model, AI system producers will incur both production and liability costs. Price and quality are the factors that differentiate systems. But we depart from the standard model

[5]Kessler and McClellan (1996) and Sloan and Shadle (2009) study defensive medicine.

[6]There are multiple explanations for this. Notably, the magnitude of the compensation payment might be insufficient to represent a significant sanction. Physicians would be fully insured against the risk of medical malpractice. Malpractice cases would not result in a higher insurance premium for the physician involved. Finally, faced with uncertainty about the legal standard of care required, physicians would be totally unaware of the effect of liability.

in order to propose a definition that is better-suited to the demand function for AI systems by physicians. In other words, we propose a closer analysis of the physician's willingness-to-pay for AI systems of varying quality. Furthermore, we replace the firms' degree of caution found in economic analyses of product liability by a quality level, based on a vertical differentiation IO model. With these changes, we consider that our model is more suited to the analysis of the effects of compensation, and its division between physicians and AI system producers in the AI healthcare market. However, unlike the standard model found in economic analyses of product liability, our model makes it possible to formally compare liability regimes (strict and fault-based). Furthermore, our model does not refer to information asymmetries, wherein the firm holds relevant information about the accuracy of the AI system that is unknown by the physician. The latter two points are not the focus of our analysis. In the rest of the paper, especially when describing the framework, we take care to be very explicit about where we depart from the standard model used to analyse product liability. We can summarise our findings as follows.

We show that at equilibrium two AI system producers each enjoy a positive market share, while some patients are still diagnosed by physicians. One producer sells one product with the best-available quality. Their competitor sells another, lower-quality system, at a lower price. The choice of quality level is independent of how any compensation paid to the patient is shared between physicians and AI system producers. However, the magnitude of the compensation does matter: in particular, the higher it is, the lower the vertical differentiation. As a consequence, lower compensation payments, as suggested by the European Commission, may degrade the quality of the low-quality AI system, while the quality of the high-quality system remains unchanged. This change in European liability law could have two opposite effects on AI system producers' profits. On the one hand, the increase in vertical differentiation reduces price competition, thus increasing firms' profits. On the other hand, demand for AI systems will decrease because physicians are even less concerned about the quality of the diagnosis given that they will have to pay less compensation if a medical error occurs. This second effect negatively impacts firms' profits.

Next, we consider that quality levels are fixed, and focus on the equilibrium price of the AI system. We show that prices increase as the compensation paid to patients in the case of a medical error increases. Furthermore, the price of the high-quality AI system rises by more than the low-quality AI system under the condition that the physician only pays a minimum level of compensation. Physicians are all the more inclined to invest in the highest-quality AI system if the compensation paid to the patient is high - under the condition that the physician pays a substantial share of the compensation to the benefit of the AI system producer. If the physician can significantly reduce her or his share of any compensation (at the expense of the AI system producer), then the highest-quality AI system becomes less attractive, and less sensitive to the magnitude of the compensation payment.

The remainder of this paper is organized as follows. Section 2 describes price competition. Section 3 describes quality competition. We discuss our results in Section 4. Section 5 concludes the paper. Appendices are given in Section 6.

2 Price competition and liability

We assume that two firms are selling AI systems of different qualities to a representative physician. These goods aid clinical decision-making and enhance the physician's judgement. The quality of this good is indicated by a variable α chosen by the firm, with higher α indicating higher quality. The variable α can be regarded as the degree of accuracy of the medical diagnosis. More precisely, we check that α equals both the probability of a true positive (sensitivity) and the probability of a true negative specificity).[7] In other words, the variable α measures the ability to correctly identify patients with or without a disease, and this quality is observable by physicians prior to purchase. Post-purchase, quality is also verifiable by a court. Finally, quality influences the probability of medical error (an incorrect patient diagnosis, as explained above), which is equivalent, in our setting, to the likelihood of patient harm. It should be noted, however, that quality does not influence the magnitude of patient harm. Below, we explore the physician's willingness-to-pay for one unit of such an AI system to diagnose a heterogeneous group of patients.

Diagnostic quality may vary across AI systems and physicians as follows. If the physician does not use an AI system, then the medical decision will be based on the physician's own capabilities. In our setting, the variable α^P measures the physician's capabilities (quality). As a consequence, the physician will treat (T) a sick patient with a probability α^P and will not treat (NT) a sick patient with a probability $(1-\alpha^P)$. Let us consider healthy patients. If the physician does not use an AI system, then she or he will not treat (NT) a healthy patient with a probability α^P, and will treat (T) a healthy patient with a probability $(1-\alpha^P)$.

[7] In our model, we do not strictly distinguish between the sensitivity and specificity of the test because we do not seek to study the extent to which each firm may simultaneously invest in plural dimensions of accuracy in order to differentiate its AI system. Here, we study what motivates firms to find an accuracy/price mix that will differentiate them from their competitors. Consequently, we do not consider the usual trade-off between sensitivity and specificity, such that higher sensitivity means lower specificity, and vice versa.

Alternatively, the physician can use an AI system to diagnose the patient. In this case, the physician will treat (T) a sick patient with a probability α^{AI} and will not treat (NT) a sick patient with a probability $1-\alpha^{AI}$ where α^{AI} measures the AI's quality, assuming $\alpha^{AI} > \alpha^{P}$. AI algorithms perform better than humans (Cheng et al. 2016). If the physician uses an AI system to diagnose the patient, then she or he will not treat (NT) a healthy patient with a probability α^{AI} and will treat (T) a healthy patient with a probability $(1-\alpha^{AI})$.

The quality of diagnosis may differ across the two competing firms. Hereafter, we denote these two AI systems or firms as 1 and 2. We consider that the firm 1 sells a low-quality AI system, while the firm 2 sells a high-quality system. With our notation, this corresponds to $\alpha_1^{AI} < \alpha_2^{AI}$. These qualities are exogenous in the price game. Furthermore, each firm can produce and sell only one type of AI system, as is usual in a duopoly market. Each physician consumes either one or zero units of AI system by patient. Physicians and AI system producers are risk-neutral. As qualities differ ($\alpha_1^{AI} < \alpha_2^{AI}$), price also differs. Using our notation, $p_1 < p_2$, where the variable p is the unit price of the AI system, that is, the per-patient price of the medical diagnosis. Finally, each firm has constant marginal costs c for each unit of AI system of quality α^{AI}. The cost of production is proportional to the quantity produced. There are no fixed costs, and marginal production costs do not depend on the quality level. As a consequence, better quality does not increase the marginal cost of production, which is standard in models of vertical product differentiation.

Demand is defined by a representative physician, who maximizes her or his utility when using and buying a high- or low-quality AI system, or even refraining from using a system to diagnose patients. To derive the physician's utility, we need to describe the benefits and costs. Let h_0 be the health status of a sick patient, and h_1 be the health status of a healthy patient, with $h_0 < h_1$. We denote b as the health benefit associated with treating a sick patient. The variable d measures the health loss in the case of wrongly treating a healthy patient.

In order to capture the effect of liability law, we introduce two variables, D and β. In the case of a wrong decision or medical error (either the decision to not treat (NT) a sick patient, or to treat (T) a healthy patient), the patient will receive a compensation payment D[8]. Both physicians and AI system producers are fully-informed about the magnitude of the compensation. We assume that this payment may be shared between the AI system producers and physicians, under the condition that the physician has used an AI

[8]For simplicity, the compensation payment is the same for both type-I and type-II errors. This simplification does not change the results of the model.

system to decide whether to treat, or not, the patient. We denote βD as the amount of compensation paid by the physician. The amount of compensation paid by the AI system producer is equal to $(1-\beta)D$. We assume that $0 < \beta < 1$. As a substantial portion of compensation may be allocated to AI system producers, the firm's marginal costs include both production and liability costs, but not litigation costs. We assume that production costs do not increase as AI quality increases, and that expected liability costs fall as the system's quality increases. Thus, it is possible for the high-quality AI system to have lower overall marginal costs than the low-quality system, and that this might reduce the price of the high-quality system. However, in our setting, AI system producers have an incentive to differentiate on quality, in order to reduce price competition. This allows high-quality AI systems to be priced higher than low-quality systems, because producers seek to avoid copying each other. Finally, one crucial point in our model is that the patient is fully compensated for harm. As a consequence, we only focus on strict liability, rather than negligence. We also do not address whether the physician or the AI producer is liable or not. In the context of AI, it could be argued that the quality standard is taken to be uncertain on an ex ante basis, that is before any medical error occurs. In this case, negligence would be formally very similar to strict liability. But, since our version of strict liability does not allow for patients to be under-compensated, we cannot consider negligence either for the physician, or for the AI system producer. In practice, under fault-based liability, patients are likely to be under-compensated, even if the quality standard is uncertain. We discuss liability regimes in more detail in the interpretation of our results (Section 4).

Benefits and costs for the physician are set as utilities U. In our model, these utilities have a three-state dependence, U_{ij}^{k}, thus leading to 2^3 combinations. The variable i denotes the patient's true state of health (S or H), j denotes the treatment decision (T or NT), and k denotes whether the physician uses AI to diagnose the patient or not (the index k equals P for the physician). For example, U_{ST}^{P} equals the physician's utility in treating a sick person based on the physician's diagnosis, while U_{HNT}^{AI} equals the physician's utility by not treating a healthy person, based on an AI system diagnosis. Using this notation, we can derive (Table 1) the physician's utility associated with each case (i, j, k). Figure 1 (below) summarizes these scenarios.

$U_{ST}^{P} = h_0 + b$	$U_{ST}^{AI} = h_0 + b - p$
$U_{SNT}^{P} = h_0 - D$	$U_{SNT}^{AI} = h_0 - p - \beta D$
$U_{HT}^{P} = h_1 - d - D$	$U_{HT}^{AI} = h_1 - d - p - \beta D$
$U_{HNT}^{P} = h_1$	$U_{HNT}^{AI} = h_1 - p$

Table 1: The physician's utility associated with each treatment decision i, j, k.

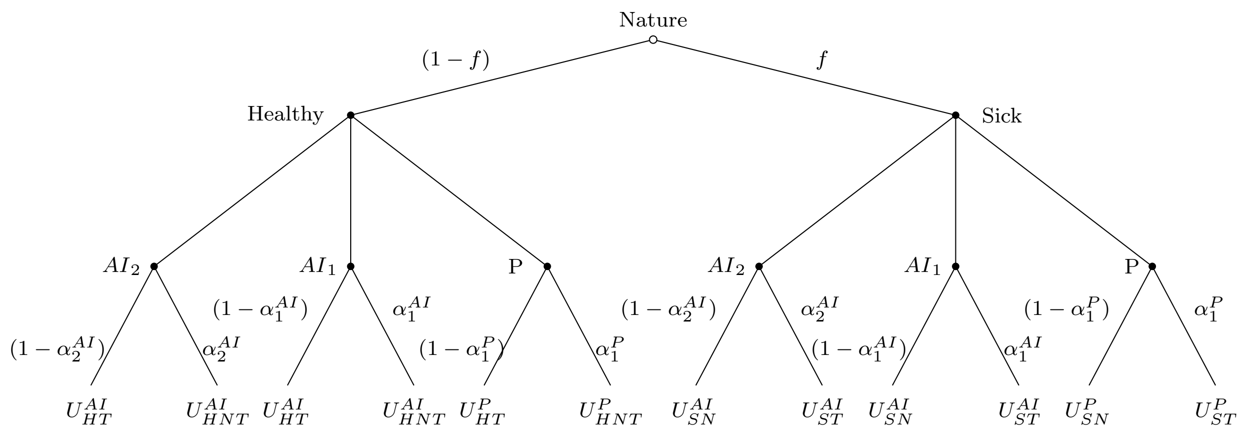

Figure 1: Probabilities in the diagnostic procedure

In Figure 1, Nature indicates the type of patient. The patient is healthy with probability $1-f$. He or she is sick with probability f. Whether the patient is healthy or sick, we focus on two ways of arriving at a correct diagnosis. First, we can treat (T) a sick patient (S). Second, we can not treat (NT) a healthy patient (H). Consequently, a medical error is the decision to either treat (T) a healthy patient (H), or to not treat (NT) a sick patient (S).

With the utilities given in Table 1, and the probabilities given in Figure 1, we can derive demand at the level of each firm. From the above, we note that the expected utility of not using an AI system is equal to:

$$f(\alpha^P(h_0+b)+(1-\alpha^P)(h_0-D))+(1-f)(\alpha^P h_1+(1-\alpha^P)(h_1-d-D)) \tag{1}$$

while, for the alternative action, the expected utility from buying an AI system produced by firm 1 or 2 (and setting it up following the producer's recommendations) is equal to:

$$f(\alpha_i^{AI}(h_0+b-p_i)+(1-\alpha_i^{AI})(h_0-p_i-\beta D))+(1-f)(\alpha_i^{AI}(h_1-p_i)+(1-\alpha_i^{AI})(h_1-p_i-d-\beta D)) \tag{2}$$

for $i=1,2$

Firm-level demand is defined as a representative physician who maximizes utility when buying a high-quality AI system rather than a low-quality system, or when refraining from buying any AI system. Given

(p_1, p_2), we use $\overline{f}(p_1, p_2)$ to denote the marginal physician who has no preference regarding the choice between buying either of the two AI systems. By definition, this marginal physician satisfies:

$$f(\alpha_1^{AI}(h_0 + b - p_1) + (1 - \alpha_1^{AI})(h_0 - p_1 - D)) + (1 - f)(\alpha_1^{AI}(h_1 - p_1) + (1 - \alpha_1^{AI})(h_1 - p_1 - d - D))$$
$$= f(\alpha_2^{AI}(h_0 + b - p_2) + (1 - \alpha_2^{AI})(h_0 - p_2 - D)) + (1 - f)(\alpha_2^{AI}(h_1 - p_2) + (1 - \alpha_2^{AI})(h_1 - p_2 - d - D)).$$

Thus $\overline{f}(p_1, p_2) = \frac{p_2 - p_1 - (\alpha_2^{AI} - \alpha_1^{AI})(d+D)}{(\alpha_2^{AI} - \alpha_1^{AI})(b-d)}$. If $f > \overline{f}(p_1, p_2)$ then the physician strictly prefers the high-quality AI-system. If $f < \overline{f}(p_1, p_2)$ then he or she strictly prefers the low-quality system.

However, the physician could decide to not use an AI system at the current price (p_1, p_2). We use $\underline{f}(p_1, p_2)$ to denote the physician who has no preference between buying a low-quality AI system or not buying an AI system at all. This marginal physician is defined as the solution to:

$$f(\alpha^P(h_0 + b) + (1 - \alpha^P)(h_0 - D)) + (1 - f)(\alpha^P h_1 + (1 - \alpha^P)(h_1 - d - D))$$
$$= f(\alpha_1^{AI}(h_0 + b - p_1) + (1 - \alpha_1^{AI})(h_0 - p_1 - D)) + (1 - f)(\alpha_1^{AI}(h_1 - p_1) + (1 - \alpha_1^{AI})(h_1 - p_1 - d - D)).$$

Thus $\underline{f}(p_1, p_2) = \frac{p_1 - (\alpha_1^{AI} - \alpha^P)(d+D)}{(\alpha_1^{AI} - \alpha^P)(b-d)}$ where $\alpha^{AI} > \alpha^P$. If $f > \underline{f}(p_1, p_2)$ then the physician will buy the low-quality AI system. If $f < \underline{f}(p_1, p_2)$ then he or she will not buy a low-quality AI system. Thus, the physician will decide to treat, or not, the patient based on their own capabilities, captured by the variable α^P. As a consequence, the market is uncovered, as some physicians will decide not to use an AI system to diagnose their patients. We consider that the decision depends on patient characteristics. The latter can either be healthy with probability $1 - f$, or sick with probability f. The variable f is uniformly distributed on the interval $(0, 1)$ from a population of N patients.

To summarise, if $f < \underline{f}(p_1, p_2)$ then the clinician does not use an AI system. If $\underline{f}(p_1, p_2) < f < \overline{f}(p_1, p_2)$ then the clinician buys the low-quality system (α_1^{AI}, p_1). If $f > \overline{f}(p_1, p_2)$ then the clinician buys the high-quality system (α_2^{AI}, p_2). Given these elements, we can derive demand functions for high-quality and low-quality AI systems.

Proposition 1 *We have a unique market configuration at price equilibrium characterized by the demand functions $D_i(p_1, p_2) > 0$ for $i = 1, 2$:*

$$D_1(p_1, p_2) = \frac{N}{b - d}\left(\frac{p_2 - p_1}{\alpha_2^{AI} - \alpha_1^{AI}} - \frac{p_1 - (1 - \alpha^P)(1 - \beta)D}{\alpha_1^{AI} - \alpha^P}\right) \tag{3}$$

for the low-quality AI system and

$$D_2(p_1,p_2) = \frac{N}{b-d}\left(b + \beta D - \frac{p_2 - p_1}{\alpha_2^{AI} - \alpha_1^{AI}}\right) \tag{4}$$

for the high-quality AI system.

Demand for both high-quality and low-quality AI systems increases with the amount of compensation paid to the patient, D. Two effects play a role in this. The first results from the improved accuracy of diagnoses due to the use of an AI system. Intuitively, as compensation D increases, physicians become more concerned about the accuracy of the AI system, because they will have to pay D in the case of a medical error. As a consequence, better diagnostic accuracy is more attractive to the physician. The second effect is that it gives the physician the option to transfer a portion of the compensation paid to the patient in the case of a medical error to the AI system producer. Thus, even a low-quality system becomes attractive, because using such a product enables physicians to transfer some of the compensation that must be paid to the producer, regardless of the quality of the system.

Next, we examine in detail the effect of the division of the compensation payment between the physician and the producer. We use β to denote the share of compensation borne by the physician, and $1-\beta$ is the compensation paid by the firm. Demand for high-quality AI systems increases as the share of compensation paid by the physician increases, while, at the same time, demand for low-quality AI systems decreases. The explanation is as follows: the more the physician has to pay in the event of a medical error, the more he or she is sensitive to the accuracy of the AI system. As demand for high-quality systems increases, demand for low-quality systems falls.

When $\beta = 0$, we also see that if two AI systems, with different levels of accuracy are equally priced, then physicians will only choose the one with the highest accuracy, as only $D_2(p_1,p_2)$ will be positive. This is a standard property of vertical differentiation.

Given the demand functions given in Equations (3) and (4), profits of the two firms in the price game are:

$$\pi_1 = (p_1 - c - (1-\beta)(1-\alpha_1^{AI})D)\left(\frac{N}{b-d}\left(\frac{p_2 - p_1}{\alpha_2^{AI} - \alpha_1^{IA}} - \frac{p_1 - (1-\alpha^P)(1-\beta)D}{\alpha_1^{AI} - \alpha^P}\right)\right) \tag{5}$$

$$\pi_2 = (p_2 - c - (1-\beta)(1-\alpha_2^{AI})D)\left(\frac{N}{b-d}\left(b + \beta D - \frac{p_2 - p_1}{\alpha_2^{AI} - \alpha_1^{AI}}\right)\right) \tag{6}$$

The respective responsibilities of the physician and the AI system producer will influence the firm's profits at two levels. First, the share of compensation borne by the producer reduces the firm's net profit margin. The first terms in Equations (5) and (6), $(1-\beta)(1-\alpha_i^{AI})D$, represent expected liability costs paid by the firm in the case of a medical error. Second, the demand functions also depend on the division of the compensation payment between the physician and the AI system producer. By differentiating Equations (5) and (6) with respect to p_1 and p_2, respectively, setting the two resulting expressions equal to zero and solving, we can establish the following proposition for equilibrium prices:

Proposition 2 *The Nash equilibrium is*

$$p_1^* = \frac{\begin{pmatrix} 3c\left(\alpha_2^{AI}-\alpha^P\right)+2b(\alpha_2^{AI}-\alpha_1^{AI})(\alpha_2^{AI}-\alpha^P)+D(1-\beta)(1-\alpha^P)(\alpha_2^{AI}-\alpha_1^{AI}) \\ +D(\alpha_2^{AI}-\alpha^P)\left(2\beta(\alpha_2^{AI}-\alpha_1^{AI})+2(1-\beta)(1-\alpha_2^{AI})+(1-\beta)(1-\alpha_1^{AI})\right) \end{pmatrix}}{4(\alpha_2^{AI}-\alpha^P)-(\alpha_1^{AI}-\alpha^P)} \tag{7}$$

$$p_2^* = \frac{\begin{pmatrix} c\left(2\left(\alpha_2^{AI}-\alpha^P\right)+(\alpha_1^{AI}-\alpha^P)\right)+b(\alpha_2^{AI}-\alpha_1^{AI})(\alpha_1^{AI}-\alpha^P) \\ +D(1-\beta)\left((\alpha_2^{AI}-\alpha_1^{AI})(1-\alpha^P)+3(\alpha_2^{AI}-\alpha^P)(1-\alpha_1^{AI})\right)+D\beta(\alpha_1^{AI}-\alpha^P)(\alpha_2^{AI}-\alpha_1^{AI}) \end{pmatrix}}{4(\alpha_2^{AI}-\alpha^P)-(\alpha_1^{AI}-\alpha^P)} \tag{8}$$

Any change in the value of the compensation paid to the patient in the case of a medical error (D) has an impact on the firm's production costs and demand. First, any increase in compensation is passed on to the consumer in the price of the AI system. Moreover, we have previously shown that demand for the AI system increases with the amount of the compensation payment. Both of these effects explain why, at equilibrium, with a given quality of AI system, an increase in the compensation payment will result in an increase in the price of both AI systems.

We show in Appendix 2 that the price of both high-quality and low-quality AI systems decreases as the share of compensation borne by the physician decreases. This result confirms that the division of compensation between physicians and AI system producers has an impact on production costs and demand. The higher the share of compensation borne by the physician, the lower the AI system's production costs. As a consequence, the price of both high-quality and low-quality systems decreases, because of price competition between producers.

Furthermore, the impact of the division of compensation between AI system producers and physicians on demand depends on the quality of the diagnosis provided by the AI system. We have previously shown that demand for the high-quality system increases with the share of compensation borne by the physician, which pushes the price of the high-quality system up. In contrast, demand for the low-quality system decreases as the share of compensation borne by the physician decreases, which pushes the price of low-quality AI systems downwards. For the low-quality AI system, these two economic effects run in the same direction, while for the high-quality system, the demand effect works in the opposite direction to the cost effect. However, the aggregate effect is clear: the price of both high-quality and low-quality AI systems decreases as the share of compensation borne by physicians decreases.

Finally, we find that an increase in the accuracy of the clinician's diagnosis (α^P) induces a decrease in the equilibrium price. This is due to competition between AI systems and clinicians, whose opinion is free. It is logical to expect that marginal costs associated with the production of the AI system (c) leads to a fall in the price of both systems. Furthermore, the greater the patient's health benefit associated with treatment, the higher the price of both systems.

2.1 Equilibrium difference in price between high-quality and low-quality AI systems

The price difference between a high-quality and a low-quality AI system equals:

$$p_2^* - p_1^* = (\alpha_2^{AI} - \alpha_1^{AI}) \left(\frac{c + (b + \beta D)((\alpha_2^{AI} - \alpha_1^{AI}) + (\alpha_2^{AI} - \alpha^P)) - 2(1-\beta)D(\alpha_2^{AI} - \alpha^P)}{4(\alpha_2^{AI} - \alpha^P) - (\alpha_1^{AI} - \alpha^P)} \right) \tag{9}$$

Clearly, AI system producers have an incentive to differentiate their products. Here, we assume that the two AI systems have identical capabilities, that is $\alpha_2^{AI} = \alpha_1^{AI}$. Thus, their prices are identical, leading to zero profit at equilibrium.

Furthermore, we noted above that the difference in price between the two AI systems increases with the share of compensation borne by the physician. We recall that both equilibrium prices decrease with the share of compensation borne by the physician. Thus, we find that the price of the low-quality AI system decreases far more than that of the high-quality system. This is because an increase in the share of compensation borne by the physician results in a corresponding increase in demand for the high-quality system, but decreases demand for the low-quality system.

We explore the impact of the magnitude of the compensation payment on the difference in price between high-quality and low-quality AI systems in Proposition 3, below.

Proposition 3 *The difference in price between high-quality and low-quality AI systems depends on the magnitude of the compensation payment, D, as follows:*

$$\frac{\partial(p_2^* - p_1^*)}{\partial D} > 0 \quad if \quad \beta > \Delta \tag{10}$$

where

$$\Delta = \frac{2(\alpha_2^{AI} - \alpha^P)}{3(\alpha_2^{AI} - \alpha^P) + (\alpha_2^{AI} - \alpha_1^{AI})} \tag{11}$$

with $0 < \Delta < 1$, $\frac{\partial \Delta}{\partial \alpha^P} < 0$, $\frac{\partial \Delta}{\partial \alpha_2^{AI}} < 0$ *and* $\frac{\partial \Delta}{\partial \alpha_1^{AI}} > 0$.

We know that at equilibrium, the price of the AI system increases with the compensation paid to patients. Here, we show that the high-quality system price increases by more than that of the low-quality system under the condition that the variable β is large enough. The latter condition stipulates that the physician pays a minimum amount of compensation. In contrast, if β is low enough, then the price of the high-quality system rises less than the price of the low-quality system. The physician will be all the more inclined to invest in the highest-quality system, under the condition that he or she pays a substantial share of the compensation. When the physician can avoid paying for medical errors, demand for the highest-quality system becomes less attractive, and less sensitive to the magnitude of the compensation paid to the patient.

2.2 Firms' net margins at equilibrium

Finally, we study firms' net margins at equilibrium. Using Equations (7) and (8), we find that:

$$\begin{aligned} p_1^* - c - (1-\beta)(1-\alpha_1^{AI})D &= \left(\frac{\alpha_2^{AI} - \alpha_1^{AI}}{4(\alpha_2^{AI} - \alpha^P) - (\alpha_1^{AI} - \alpha^P)}\right)\left(-2c + (b+D)(\alpha_1^{AI} - \alpha^P)\right). \\ p_2^* - c - (1-\beta)(1-\alpha_2^{AI})D &= \left(\frac{\alpha_2^{AI} - \alpha_1^{AI}}{4(\alpha_2^{AI} - \alpha^P) - (\alpha 1^{AI} - \alpha^P)}\right)\left(-c + 2(b+D)(\alpha_2^{AI} - \alpha^P)\right). \end{aligned}$$

Here, the main result is that at equilibrium net margins are independent of the share of compensation borne by the physician (β). Regardless of the division of compensation between the physician and the producer, any increase in liability costs paid by the producer to the benefit of the physician will be reflected in the price of the AI system, at the expense of physicians. Intuitively, the amount of the compensation payment that is transferred from the physician to the producer is inevitably passed on to the consumer in the price of the AI system.

3 Liability and quality choices

In this section, we study the choice of quality levels by firms in an interval $(\alpha^P, 1)$. Given the demand functions set by equations (3) and (4) and the equilibrium price given by Equations (7) and (8), we can calculate the two firms' profit functions as a function of the quality of the two competing AI systems:

$$\pi_1 = \left(\frac{N}{b-d}\right)(\alpha_2^{AI} - \alpha_1^{AI})\left(\frac{\alpha_2^{AI} - \alpha^P}{\alpha_1^{IA} - \alpha^P}\right)\left(\frac{-2c + (b+D)(\alpha_1^{AI} - \alpha^P)}{4(\alpha_2^{AI} - \alpha^P) - (\alpha_1^{AI} - \alpha^P)}\right)^2. \tag{12}$$

$$\pi_2 = \left(\frac{N}{b-d}\right)(\alpha_2^{AI} - \alpha_1^{AI})\left(\frac{-c + 2(b+D)(\alpha_2^{AI} - \alpha^P)}{4(\alpha_2^{AI} - \alpha^P) - (\alpha_1^{AI} - \alpha^P)}\right)^2. \tag{13}$$

First, we consider the high-quality firm that corresponds best to α_1^{IA}. We show in Appendix 2 that the profit of the firm that produces the high-quality AI system, defined by Equation (13), is strictly increasing as quality improves, that is, with the variable α_2^{AI}. Consequently, at the second stage of equilibrium, the firm producing the high-quality AI system will set the quality level as high as possible, $\alpha_2^{AI^*} = 1$. There are a number of reasons for this. First, demand for the AI system increases with quality. Second, for a given quality level of the other AI producer, price competition is reduced. Finally, better quality draws more physicians into the market for AI systems.

Next, we develop the analysis of the producer of the low-quality AI system that corresponds to $\alpha_2^{AI^*} = 1$. Choosing $\alpha_1^{AI} = \alpha_2^{AI} = 1$ cannot be optimal because it results in Bertrand competition in the price game. Given Equations (12) and (13), the firm's profits would be equal to zero. When choosing the quality level, the producer of the low-quality AI system will solve $\frac{\partial \pi_1}{\partial \alpha_1^{AI}} = 0$, given that the firm producing the high-quality AI system will choose $\alpha_2^{AI^*} = 1$. Solving this equation, we find that the high-quality producer will set quality to the highest possible value, while its competitor will set quality to a strictly lower value, $\alpha_1^{AI^*} < \alpha_2^{AI^*} = 1$. This equilibrium value $\alpha_1^{AI^*}$ is implicitly defined by:

$$\frac{-2c + (b+D)(\alpha_1^{AI^*} - \alpha^P)}{-c + 2(b+D)(1-\alpha^P)} - 4\left(\frac{1-\alpha_1^{AI^*}}{4(1-\alpha^P) - (\alpha_1^{AI^*} - \alpha^P)}\right)\left(\frac{\alpha_1^{AI^*} - \alpha^P}{1-\alpha^P}\right) = 0 \tag{14}$$

Using the Implicit Function Theorem, we can check that the optimal level for the low quality producer increases with increasing compensation, D, while it is independent of the share of the compensation payment borne by the firm, β. As a consequence, the higher the liability cost, the higher the low-quality level. Thus, any increase in the compensation paid to the patient in the case of medical malpractice will decrease vertical differentiation, thereby increasing price competition between AI system producers. Our main findings are summarised in the next proposition:

Proposition 4 *The two AI system producers enjoy positive market shares, while some patients are still diagnosed by physicians without additional support from an AI system. One firm produces at the best-available quality. The other sells an AI system with strictly lower quality, at a lower price. The choice of quality is independent of how the compensation paid to the patient is divided between physicians and producers. But the magnitude of this compensation does matter. The higher the payment, the lower the vertical differentiation. This has two, opposing effects on firms' profits. On the one hand, the reduction in vertical differentiation increases price competition, thus reducing firms' profits. On the other hand, demand for AI systems increases with the compensation paid to the patient because physicians are all the more concerned about the quality of the diagnosis if they will have to pay more compensation in the case of a medical error. This second effect positively impacts firms' profits.*

4 Discussion

4.1 Liability rules

We propose a general scheme in which the patient is paid D in compensation. This compensation may be divided between the physician and the AI system producer, β being the share borne by the physician. Regarding liability regimes, this scheme is consistent with strict liability rules. Under strict liability, the firm or the physician will be liable without the patient having to prove negligence or fault. When harm is caused, the court holds both the physician and the firm liable, and orders both parties to pay compensation to the patient. In practice, it is often the case that both defendants pay an equal share, such that $\beta = 1/2$.

Strict liability can also apply to defective product lawsuits. In this case, the AI system producer may be found liable, even if the patient has no proof of the defendant's negligence or intent to harm. It is simply necessary to prove the AI system is defective. However, this is not the case for the physician, where it is usually necessary to prove fault or negligence. It could be said that there is a regulatory mix between fault-based and strict liability regimes. The physician who fails to make an accurate diagnosis faces fault-based liability in tort, should that decision injure the patient. The AI system producer faces strict liability for defective medical products. Thus the producer can be held liable for harm caused by a defective AI system.

If the physician does not use an AI system to diagnose the patient, and makes a wrong diagnosis, he or she is considered to not meet the standard of care. As this conduct is negligent, the patient must be compensated. If the physician does use an AI system to diagnose the patient, the standard of care is met, but we can assume that, with a probability β, there was a technical error when using the system. In such a case, the physician is at fault and pays compensation. As for the complementary probability $1 - \beta$, the physician cannot be held liable for any failure to diagnose the patient. In this case, the patient can make a compensation claim against the AI system producer under strict liability law, and the firm is responsible for any payment to the patient.

4.2 Healthcare quality and the AI system producer's liability

Product liability rules aim to maintain a fair balance between the interests of users and producers[9]. Producer protection is often understood as trying to limit their liability and the possible amounts of compensation paid to consumers if a product is defective. One of the main arguments for limiting producers' liability is that transferring extremely high costs to firms would limit their ability to innovate, or even propose innovative products due to increased production costs. We show in this paper that the most important effect is the magnitude of the compensation D paid either by the firm or the physician, following a misdiagnosis. The magnitude of this compensation can be understood as the degree of malpractice pressure on healthcare market actors (firms and physicians) chosen by the policymaker. Even if, in practice, governments and the insurance sector bear some of this liability cost, we show that an increase in malpractice pressure increases demand for AI systems, leading to an increase in *defensive medicine*. Furthermore, under malpractice pressure, the average quality of AI systems improves. As product differentiation diminishes, increased competition between firms reduces prices. As a consequence, the effect of malpractice pressure on both prices and profits is ambiguous. Higher demand raises prices, but reduced product differentiation lowers them. Finally, increased malpractice pressure will be beneficial for patients as healthcare quality improves, with mixed effects on prices.

Next, we look at how our results depend on the extent to which malpractice pressure is oriented towards either firms or physicians. Malpractice pressure is captured by the variable β, that is, the division of compensation between physicians and AI system producers. We show that this element has no impact on quality and profits at equilibrium. If there is a change in the division of compensation between producers and

[9]Directives regarding the liability associated with defective products can be found on the European Commission's website https://ec.europa.eu/growth/single-market/goods/free-movement-sectors/liability-defective-products_en.

physicians, firms simply transfer this change into prices. At the same time, demand for AI systems remains unchanged, because the reduction in physicians' liability associated with the use of systems will be cancelled by higher prices for these products.

5 Conclusion

To conclude, we briefly list some related research questions that may be important in the future, and that go beyond the economic analysis of liability. First, the adoption of AI systems may influence physicians' skills (assumed to be exogenous in our model). If the applications currently in use that are studied here are replaced by more intelligent and independent machines that surpass human capabilities, it is plausible that physicians may start losing confidence in their own judgement. It could also reduce their capabilities. However, AI systems and physicians could become complementary to each other, rather than substitutes. For instance, AI systems currently use data generated by physicians. Thus, we need to explore how physicians will behave if they become a fundamental part of AI algorithms. It might even be a condition for society to start trusting AI, before AI systems become independent decision-makers. Finally, the adoption of AI algorithms may be seen as a threat to clinicians. If employers begin to bargain on the rent from AI systems gains, they may use them as a way to increase their bargaining power. This could have negative effects on physicians' salaries or their role in the healthcare system.

References

[1] Acemoglu, D., 2022. Harms of AI, in *The Oxford Handbook of AI Governance*, edited by J. Bullock, Y-C. Chen, J. Himmelreich, V.M. Hudson, A. Korinek, M. Young, and B. Zhang.

[2] Aggarva, A. Matheny M.E., Shachar C., Wang S.X.Y. Thadaney-Israni S., 2022, Artificial Intelligence in Healthcare, in *The Oxford Handbook of AI Governance*, edited by J. Bullock, Y-C. Chen, J. Himmelreich, V.M. Hudson, A. Korinek, M. Young, and B. Zhang.

[3] Avraham, R., 2007. An Empirical Study of the Impact of Tort Reforms on Medical Malpractice Settlement Payments. *The Journal of Legal Studies.* 36(2), 183–229

[4] Bertoli, P., and Grembi, V., 2018. Medical Malpractice: how Legal Liability Affects Medical Decisions, in B.H. Balgati and F. Moscone (eds), *Health Econometrics*, Bingley, Emerald Publishing. 235-261

[5] Cheng, J-Z., Ni, D., Chou, Y-H., Qin, J., Tiu, C-M., Chang, Y-C., Huang, C-S., Shen, D., Chen, C-M., 2016. Computer-Aided Diagnosis with Deep Learning Architecture: Applications to Breast Lesions in US Images and Pulmonary Nodules in CT Scans. *Scientific Reports.* 6(1), 1–13

[6] Danzon, P.M., 1985a. Liability and Liability Insurance for Medical Malpractice. *The Journal of Health Economics.* 4, 309–331

[7] Danzon, P.M., 1985b. Medical Malpractice: Theory, Evidence, and Public Policy, Cambridge (Mass.), Harvard University Press.

[8] Danzon, P.M., and Lillard, L.A., 1982. Settlement Out of Court: the Disposition of Medical Malpractice Claims. *The Journal of Legal Studies.* 12(2), 345–377

[9] Daughety, A., and Reinganum, J., 2013. Economic Analysis of Products Liabilty: Theory. In *Research Handbook on the Economics of Torts.* Edward Elgar Publishing.

[10] De Chiara, A., Elizalde, I., Manna, E., 2021. Car Accident in the Age of Robots. *The International Review of Law and Economics.* 68

[11] Farber, H.S., and White, M.J., 1991. Medical Malpractice: an Empirical Examination of the Litigation Process. *The Rand Journal of Economics.* 22(2), 199–217

[12] Frakes, M., and Jena, A.B., 2016. Does Medical Malpractice Law Improve Health Care Quality. *Journal of Public Economics.* 143, 142–158

[13] Gaynor, M., Ho, K., Town, and Robert J., 2015. The Industrial Organization of Health-Care Markets. *The Journal of Economic Literature.* 53(2), 235–284

[14] Guerra A., Parisi, F., Pi, D., 2021. Liability for Robots: an Economic Analysis. Available at SSRN 39399486.

[15] Higgins B.W. 2022, Legal Elements of an AI Regulatory Permit Program, in *The Oxford Handbook of AI Governance*, edited by J. Bullock, Y-C. Chen, J. Himmelreich, V.M. Hudson, A. Korinek, M. Young, and B. Zhang.

[16] Iizuka, T., 2013. Does Higher Malpractice Pressure Deter Medical Errors? *Journal of Law and Economics.* 56, 161–188

[17] Kessler, D., and McClellan, M., 1996. Do Doctors Practice Defensive Medicine. *The Quaterly Journal of Economics.* 111, 353–390

[18] Kim, B., 2007. The Impact of Malpractice Risk on the Use of Obstetrics Procedures, *The Journal of Legal Studies.* 36(2), 121–142

[19] Shavell, S., 2020. On the Redesign of Accident Liability for the World of Autonomous Vehicles. *The Journal of Legal Studies.* 49(2), 243–285

[20] Sloan, F., and Shadle, J.H., 2009. Is there Empirical Evidence for Defensive Medicine? a Reassessment. *Journal of Health Economics.* 28(2), 481–491

6 Appendices

6.1 Appendix 1

By differentiating Equation (11) with respect to β, we find that:

$$\frac{\partial p_1^*}{\partial \beta} = \frac{D(\alpha_2^{IA} - \alpha^P)\left[2(\alpha_2^{IA} - \alpha_1^{IA}) - 2(1-\alpha_2^{IA}) - (1-\alpha_1^{IA})\right] - D(1-\alpha^P)(\alpha_2^{IA} - \alpha_1^{IA})}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}$$

Rearranging this equation gives:

$$\frac{\partial p_1^*}{\partial \beta} = \frac{-D(1-\alpha_2^{IA})\left[2(\alpha_2^{IA} - \alpha^P) + (\alpha_2^{IA} - \alpha^P) + (\alpha_2^{IA} - \alpha_1^{IA})\right]}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)} < 0.$$

By differentiating Equation (12) with respect to β, we find that:

$$\frac{\partial p_2^*}{\partial \beta} = \frac{-D\left[(\alpha_2^{IA} - \alpha_1^{IA})(1-\alpha^P) + 3(\alpha_2^{IA} - \alpha^P)(1-\alpha_1^{IA})\right] + D(\alpha_1^{IA} - \alpha^P)(\alpha_2^{IA} - \alpha_1^{IA})}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}$$

Rearranging this equation gives:

$$\frac{\partial p_2^*}{\partial \beta} = \frac{-D(1-\alpha_1^{IA})\left[(\alpha_2^{IA} - \alpha_1^{IA}) + 3(\alpha_2^{IA} - \alpha^P)\right]}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)} < 0.$$

6.2 Appendix 2

$$\pi_2 = \left(\frac{N}{b-d}\right)(\alpha_2^{IA} - \alpha_1^{IA})\left(\frac{-c + 2(b+D)(\alpha_2^{IA} - \alpha^P)}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}\right)^2$$

The high-accuracy AI firm's decision. Deriving profit relative to quality α_2^{IA}, we have:

$$\frac{\partial \pi_2}{\partial \alpha_2^{IA}} > 0 \Leftrightarrow \frac{-c + 2(b+D)(\alpha_2^{IA} - \alpha^P)}{-2c + (b+D)(\alpha_1^{IA} - \alpha^P)} > 4\left(\frac{\alpha_2^{IA} - \alpha_1^{IA}}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}\right)$$

Substituting with the equilibrium price given in Equation (39), we have:

$$\frac{\partial \pi_2}{\partial \alpha_2^{IA}} > 0 \Leftrightarrow \frac{p_2 - c}{p_1 - c} > 4\left(\frac{\alpha_2^{IA} - \alpha_1^{IA}}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}\right)$$

As $p_2 > p_1$ at equilibrium, we have $\frac{p_2-c}{p_1-c} > 1$. Next

$$4\left(\frac{\alpha_2^{IA} - \alpha_1^{IA}}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}\right) < 1 \Leftrightarrow 3(\alpha_1^{IA} - \alpha^P) > 0.$$

Consequently

$$\frac{p_2 - c}{p_1 - c} > 1 > 4\left(\frac{\alpha_2^{IA} - \alpha_1^{IA}}{4(\alpha_2^{IA} - \alpha^P) - (\alpha_1^{IA} - \alpha^P)}\right)$$

or

$$\frac{\partial \pi_2}{\partial \alpha_2^{IA}} > 0$$