

Human Responses to AI Oversight: Evidence from Centre Court*

David Almog[†], Romain Gauriot[‡], Lionel Page[§], and Daniel Martin[¶]

April 10, 2026

Abstract

We provide field evidence on how human decision-makers respond when AI is used to audit and overrule their decisions, which we call “AI oversight”. Specifically, we study the introduction of Hawk-Eye review for close line calls in professional tennis, a high-stakes setting with objectively correct answers. We find that umpires lowered their overall mistake rate after the introduction of Hawk-Eye review, but also that umpires increased the rate at which they called balls in, producing a shift from making Type II errors (calling a ball out when in) to Type I errors (calling a ball in when out). We structurally estimate the impact of AI oversight using a model of attention-constrained umpires, and our results suggest that AI oversight increased the perceived cost of Type II errors by 37% on decisions that were overturned.

*We thank Lucas Lippman, Anastasiia Morozova, and Leshan Xu for their excellent research assistance. We are very grateful to the executives at Hawk-Eye Innovations Ltd and a leading tennis official at the time of Hawk-Eye’s introduction for providing valuable institutional information. We also thank David Argente, Meghan Busse, Colin Camerer, Andrew Caplin, Ben Enke, George Georgiadis, Mitchell Hoffman, Alex Imas, John List, Christopher Mills, Devin Pope, Michael Powell, Ashesh Rambachan, Antonio Rangel, Paul Raschky, Raphaël Raux, Luis Rayo, Yuval Salant, Jörg Spenkuch, Andrew Wooders, and audiences at EC 2024, BDRM, AI and the Future of Work at Wharton, ESIF Economics and AI+ML Meeting, Machine Learning in Economics Summer Conference, Zurich Workshop in AI+Economics, WEBEAS, MIT Sloan, Monash, UT Dallas, UH Manoa, Kellogg Strategy Brown Bag Seminar, and Booth Behavioral Lab for many helpful comments. We thank the Sloan Foundation for supporting this research under the “Cognitive Economics at Work” grant.

[†]Kellogg School of Management, Northwestern University, david.almog@kellogg.northwestern.edu.

[‡]Deakin University, romain.gauriot@deakin.edu.au.

[§]University of Queensland, lionel.page@uq.edu.au.

[¶]University of California, Santa Barbara, danielmartin@ucsb.edu.

1 Introduction

In the coming years, firms will have the option to use artificial intelligence (AI) systems to audit and overrule worker decisions across a wide array of settings. This is due to the confluence of two forces. First, machine learning algorithms are becoming increasingly good at predicting human mistakes. This includes important decisions such as bail judges predicting pretrial misconduct (Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan 2018), radiologists predicting pneumonia from chest X-rays (Rajpurkar, Irvin, Zhu, Yang, Mehta, Duan, Ding, Bagul, Ball, Langlotz, et al. 2017; Topol 2019), and workforce professionals predicting productivity for hiring and promotion (Chalfin, Danieli, Hillis, Jelveh, Luca, Jens, and Mullainathan 2016). Second, there has been a substantial drop in the cost of monitoring worker behavior, brought about by the rise in digitization.¹

We refer to a process that uses AI to audit and overrule decisions as “AI oversight”. For a concrete example of AI oversight, consider Zalando, a leading e-commerce company specializing in fashion retail. Zalando allows category managers to suggest discounts based on their private information about fashion trends. Huelden, Jascisens, Roemheld, and Werner (2024) show that while human interventions appear promising, the negative impact of poor interventions (stemming from overconfidence and other biases) offsets the benefits of good interventions (based on private information). However, they show that using algorithmic tools to predict and block undesirable human interventions has the potential to recoup these gains. In fact, recent papers in the economics literature have shown that there are large potential gains from using AI to overrule human mistakes in other high-stakes settings, including law (Rambachan 2024) and medicine (Raghu, Blumer, Corrado, Kleinberg, Obermeyer, and Mullainathan 2019).

Using AI to overrule human errors appears to be a straightforward improvement for societal welfare. However, assessing the full impact of AI oversight requires understanding whether its presence alters human decision-making. While a large amount of research has focused on how humans respond when assisted by AI, very little is known about how humans respond when their decisions might be overruled by AI. Because being overruled can carry psychological costs (e.g., shame and embarrassment of being overruled) and psychological benefits (e.g., relief at having their mistakes fixed), individuals might alter their decision-making under AI oversight.²

However, determining whether AI oversight impacts human decision-making is challenging for several reasons. First, it is necessary to have observations on decision-making both

¹The strong impact of these complementary forces (digitization and AI) is demonstrated by L. Raymond (2024) in the real estate market.

²This illustrates how insights from behavioral economics can play a key role in understanding AI-human interactions (Camerer 2019).

with and without AI oversight, but firms may be resistant to share such data given the legal and public relations risks involved. Second, to determine the types of errors people make, we need a clear “ground truth” (an objectively correct answer). In some cases, the closest one can get is by aggregating many expert opinions, such as in diagnostic radiology (Agarwal, Moehring, Rajpurkar, and Salz 2023). Third, to get a full picture of the types of errors people make, we need to observe or estimate counterfactual outcomes, as not to run afoul of incomplete data due to “selective labels” (Rambachan 2024). Fourth, it is necessary to control or otherwise account for the private information of decision-makers, as this may introduce confounds in analysis.

Given these challenges, we identified the introduction of AI oversight in umpiring at top tennis tournaments as a promising field setting for studying human responses to AI oversight. AI oversight was introduced in top tennis tournaments through an AI technology known as Hawk-Eye. Due to the very small distances involved, camera imaging alone is not sufficient to determine whether a ball lands in or out of bounds, so Hawk-Eye leverages state-of-the-art machine learning algorithms to use the ball’s trajectory and spin to reconstruct the landing point.³ Starting in 2006, players were given the option to challenge a line call, and if the call contradicted the highly accurate predictions of Hawk-Eye, the umpire’s decision was overturned, or in some cases, the point was replayed. While this form of AI oversight is initiated by humans instead of AI, it still involves using AI to audit and overturn human decisions.

By combining our proprietary Hawk-Eye data with publicly available data on player challenges, we can address many of the empirical challenges in studying the human response to AI oversight. Because we have Hawk-Eye data from the period immediately before the introduction of AI review, we can make a direct comparison before and after the introduction of AI review. This comparison is especially clean because many important factors stayed the same over the period we study: there were no substantial changes to the training and performance assessments of umpires, the pool of umpires, or the positioning and instructions to line judges.⁴ Also, because Hawk-Eye is nearly perfect, we can side-step common problems in studying human and AI interaction: uncertainty about ground truth, non-observability of counterfactual outcomes, and agents’ private information.

In this setting, we uncover strong evidence that human decision-making changes under AI oversight. We find that after the introduction of AI review, umpires lowered their overall mistake rate by 8% (1.1 p.p.) for close calls (balls landing within 100 mm of the line). However, this aggregate decrease masks two important sources of heterogeneity: the distance of the ball from the line and whether the ball was in or out. For balls just outside of the

³For similar reasons, Hawk-Eye review is being considered for sports like American football and soccer, where camera imaging alone can be insufficient to provide a definitive assessment.

⁴These institutional details were verified by a leading tennis official at the time of Hawk-Eye’s introduction.

line (within 20 mm), the mistake rate actually *increased* by 34% (8.5 p.p.). This puzzling result is explained by a behavioral response that occurs with AI oversight. We find that for the closest calls, umpires increased the rate at which they called balls in by 12.6% (6.2 p.p.) after the introduction of Hawk-Eye, which produced a shift from making Type II errors (calling a ball out when in) to Type I errors (calling a ball in when out).

This behavioral reaction can be rationalized with asymmetric costs of AI oversight. For umpires, improperly stopping a point (calling a ball out when it was actually in) became a new cause of concern with the introduction of Hawk-Eye. Even if a player successfully challenges such a call, there is no way to resume the point where it stopped, so the rules dictate that the umpire must decide whether to replay the point or award it to the challenger. Thus, Type II errors (calling a ball out when in) carry two costs: one from the impossibility of implementing the correct outcome (continuing the point from the point it was stopped) and another from having to implement an arbitrary decision that, in many cases, unleashes the outrage of the players involved and the audience. The ATP Tour did not use Hawk-Eye to assess umpires during the period that we study, suggesting that these costs are more psychological than directly related to career concerns.

While our reduced-form estimates suggest that the perceived costs of being overruled by AI led umpires to shift the rate of Type I and II errors they made, they do not indicate the magnitude of these costs. To get a sense for this, we structurally estimate a lower bound on the perceived costs of being overruled by AI using a model of attention-constrained umpires. In this model, putting in more attention is effortful (cognitively costly) but results in more accurate predictions. The umpire trades off these cognitive costs with the perceived costs of incorrect calls and being overruled by AI. The umpire has two levers for achieving this balance, and we find supportive evidence of both channels: the umpire can increase their attentional effort and vary the threshold belief at which they call a ball in or out, shifting the fraction of Type I and Type II errors.

We employ a two-stage approach to estimate the parameters of this model. We first use decisions before Hawk-Eye was introduced to recover the perceptual costs of making correct calls, and we then use decisions after Hawk-Eye was introduced to determine the perceived costs of being overruled by AI. The resulting estimates suggest that AI oversight increased the perceived cost of Type II errors by 37% on decisions that were overturned.

Two possible concerns about the external validity of these results for other professional settings are that sports officiating is a highly visible activity and Hawk-Eye is nearly perfect.⁵ First, because sports officiating is a highly visible activity, our results are more likely to extend to contexts where the outcome and overruling mechanism are observed by an audience

⁵In Appendix B, we provide a more detailed assessment of the external validity (EV) of this setting using the approach proposed in List (2020).

that the decision-maker would care about. For example, the decision-maker’s audience could be their customers, co-workers, or the general public. Second, because Hawk-Eye is nearly perfect,⁶ our results are more likely to extend to settings where AI is highly accurate *when used for corrections*. For example, while fully autonomous self-driving cars have not yet taken over the streets, many vehicles already include lane-keeping assistance systems that intervene when a driver drifts. These systems, though imperfect, operate with a high degree of confidence in the moments they are engaged, similar to other decision-support technologies that are used in professional settings.

Despite these potential limitations, our results and theoretical framework provide broader lessons on the impact of AI oversight. While introducing an AI system to overrule apparent human mistakes seems promising in principle, especially since it might motivate higher effort from decision-makers, there is a strong reason to approach this with caution: incentive misalignment. The particular implementation guidelines of the oversight system can shift the incentives for the decision-maker. This becomes problematic when the incentive shift reduces the overall quality of their decisions from the perspective of the social planner. As we argue in Section 4, we do not have reason to believe that before the introduction of Hawk-Eye review a particular type of incorrect call was costlier than the other, at least from the perspective of the umpire. Under this assumption, umpires should be aiming to minimize the total number of mistakes when there was no oversight. However, if the umpires’ attention has already reached a point where improvement would demand a tremendous amount of additional effort, distorting the relative costs of the two types of errors could result in an increase in total number of mistakes. Umpires may be inclined to reduce the now costlier type of error (e.g., calling a ball out when it was in) at the expense of increasing, by more than a 1:1 ratio, the now relatively less costly type of error (e.g., calling a ball in when it was out). Our analysis suggests that when balls landed very close to the line, umpires were more inclined to call them in, as an attempt to minimize the occurrence of the more costly type of error. It also suggests that, in overall terms, the effect of incentive misalignment outweighed the incentives to improve performance, which led to an upsurge of incorrect calls during balls that were just out. This shift in the types of errors can be highly detrimental in medical and judicial contexts, where Type I and Type II errors have widely differing implications.

For managers, the implication is that deploying oversight is a redesign of incentives and workflows, not a plug-in quality fix. Also, because oversight can improve average performance while shifting decision thresholds and reallocating mistakes across error types, firms should track error composition (not just aggregate accuracy), anticipate state-dependent responses near the boundary of human capability, and pay particular attention to the symmetry and visibility of the correction process (e.g., whether the correct outcome can be implemented

⁶A potentially interesting follow-up question is whether human reactions to AI oversight differ by the ability of the AI to overrule human decisions correctly.

cleanly, whether overrides require discretionary remedies, and whether being overruled is publicly salient). Where these features create asymmetric social or operational costs, effective deployment will typically require complementary changes to escalation rules, training, feedback, and automation of the hardest edge cases.

In this regard, our study also provides additional perspective to the debate surrounding when formal decision-making authority should be given to humans versus AI (Ide and Talamas 2025; Athey, Bryan, and Gans 2020). This is nearly certain to be an issue that policymakers, regulators, and legal authorities are forced to grapple with in the coming years.

The economics literature on AI and human decision-making, which we review in Section 1.1, has focused almost exclusively on the paradigm in which humans are provided AI recommendations. One reason for focusing on this paradigm is that giving humans final decision rights is more socially acceptable than handing these rights over to AI. We take this literature in a new direction by considering a paradigm where AI systems are used to overrule human mistakes. Despite the negative social desirability of this paradigm, firms are likely to consider using AI to overrule human mistakes because AI recommendations can fail to improve choices — for example, when decision makers hold biased beliefs (Agarwal, Moehring, Rajpurkar, and Salz 2023; Caplin, Deming, S. Li, Martin, Marx, Weidmann, and Ye 2025; Kovach, Martin, and Tserenjigmid 2026). On the other hand, AI oversight appears to improve aggregate mistake rates in our setting. Before the introduction of Hawk-Eye review, points were incorrectly awarded 13.9% of the time for close calls. After the introduction of Hawk-Eye review, points were incorrectly awarded 9% of the time. 28.6% of this 4.9 pp mistake reduction is due to improvements in umpire performance, and the remaining 71.4% is due to AI overruling incorrect human decisions.⁷

The rest of the paper proceeds as follows. In Section 2, we provide additional details on tennis umpiring, Hawk-Eye review, and the data sources we leverage in our analysis. In Section 3, we examine overall mistake rates, types of calls, and types of errors. In Section 4, we provide a model of perceptually-constrained umpires and use it to structurally estimate the psychological costs of AI oversight. In Section 5, we explore additional sources of heterogeneity. Finally, we conclude with a brief discussion in Section 6.

1.1 Related Literature

While AI is increasingly good at prediction, humans are kept in the decision process (kept “in the loop”) because of social reasons (tradition, comfort, fairness, etc.), due to labor market concerns (union power, contractual responsibilities, inequality considerations, etc.), because

⁷Given this improvement, it may be unsurprising that Major League Baseball (MLB) has followed the ATP in implementing Hawk-Eye review.

of the costs of implementing AI systems, and because humans can sometimes perform better by incorporating additional information, understanding context, handling edge cases, and adapting to changing circumstances. One form of joint AI and human decision-making that has been actively studied is when AI provides recommendations to human decision-makers. Grimon and Mills (2025) show that Child Protective Services workers are able to reduce child hospitalizations when provided with an algorithmic risk score. However, unexpectedly, the gains did not seem to come from following the algorithmic predictions but rather from reallocating their attention to different features of the allegations. L. Raymond (2024) finds that the availability of algorithmic tools in the housing market generates responses from non-adopter investors as well, shifting their focus to market sectors where human investors have comparative advantages over algorithms. Cho (2023) uses a quasi-experimental sports setting to examine how baseball umpires are affected when they work with AI assistance, focusing on their skill development. Kreitmeir and Raschky (2024) use a ban on ChatGPT in Italy to show that high-productivity programmers have worse output in the presence of AI assistance.⁸

Making decisions in environments with algorithmic recommendations elevates the need for humans to exercise proper discretion, which can lower the effectiveness of these recommendations. The findings from Hoffman, Kahn, and D. Li (2018) highlight challenges in exercising discretion in hiring, as managers are observed overruling recommendations due to personal biases rather than private information motives. Relatedly, Kawaguchi (2020) reports field experimental evidence from a retail business, showing that workers' adherence to algorithmic recommendations varies systematically with contextual factors, including sales volatility and worker experience. Exploring how humans adopt algorithmic recommendations, Agarwal, Moehring, Rajpurkar, and Salz (2023) find that providing radiologists with access to AI predictions does not, on average, result in improved performance. Another noteworthy finding from their study is that radiologists take significantly more time to reach a decision when AI information is provided.

Because giving AI final decision rights is potentially problematic, why would firms not simply provide workers with AI recommendations instead? The results above highlight two potential reasons why. First, humans may struggle to exercise appropriate discretion, sometimes deferring to AI when they should rely on their own judgment and sometimes overriding it when it is superior. For example, AI guidance can be systematically ignored due to overconfidence (Caplin, Deming, S. Li, Martin, Marx, Weidmann, and Ye 2025), and decision makers may apply an overly constant weight to algorithmic recommendations even when the informativeness of their private information varies substantially (Balakrishnan,

⁸One potential advantage of our setting for studying human-AI interaction is that, like the baseball setting, human line judges in tennis do not have any clear source of private information. We thank Ashesh Rambachan for raising this point.

Ferreira, and Tong 2026), leading to predictable over- and under-reliance. Second, in settings where a timely final decision is required, the very process of interpreting and integrating recommendations can create operational delay.

With AI oversight, discretion, a prominent force in algorithmic recommendation systems, no longer plays a major role. However, it introduces new behavioral forces into play, such as shame, pride, embarrassment, and stress. Because of this connection, we contribute to the literature on social image and peer effects (Bursztyrn and Jensen 2017) by showing that the cost of being corrected in public can overpower the potential incentive to free ride on technology. Related field evidence underscores the importance of observability in human-AI collaboration. In a recommendation-based setting, Almog (2025) finds that making workers’ reliance on AI visible to an evaluator reduces adoption and worsens performance, consistent with concerns that high reliance on AI is interpreted as a negative signal of competence. In our work, we estimate a lower bound for the psychological costs of being overruled by AI, and shame and embarrassment are potential components of these costs.⁹ A potentially revealing feature of our empirical context is that, with the introduction of AI oversight, one type of mistake became more controversial and led to a larger backlash, including complaints from players and fans. This introduced a change in the costs of different error types, which led to a decrease in the number of controversial mistakes at the expense of an increase in the number of less embarrassing ones. It is worth noting that Hawk-Eye decisions were available for television broadcasts both before and after the introduction of AI review, so this aspect of embarrassment does not change over the period we study.

We also add to the literature on monitoring, which has studied human reactions to being monitored across various domains, including marketing contexts involving online consumption habits and advertising intrusiveness (John, Acquisti, and Loewenstein 2011; Goldfarb and Tucker 2011), auditing and corruption (Olken 2007; Bobonis, Cámara Fuertes, and Schwabe 2016), and workforce productivity (Gosnell, List, and Metcalfe 2020). Nagin, Rebitzer, Sanders, and Taylor (2002) find that many call center workers behave as in the “rational cheater” model, which predicts that shirking behavior will decrease when monitoring increases. In line with this, we show that using AI systems to increase monitoring can improve performance, even when introducing new incentives to shirk by free-riding on AI corrections. Relatedly, Margolin, Reimer, and Schaupp (2024) study the introduction of the Video Assistant Referee (VAR) in German professional soccer, a system of real-time post-decision feedback that allows initially wrong calls to be reversed. They show that immediate review changes referees’ effort and decision-making, but the effects are heterogeneous: experienced referees improve final decisions without harming initial decisions, whereas inexperienced referees make worse initial decisions and do not improve final decisions. Their setting differs

⁹Our approach adds to existing methods for quantifying the welfare effects of shame and pride (Butera, Metcalfe, Morrison, and Taubinsky 2022).

from ours, however, because review is conducted by a human video official, supported by computer-aided analysis, rather than by a high-accuracy AI system that audits and overrules decisions.

Recent work on human–algorithm interaction (Caro, Colliard, Katok, Ockenfels, Stier-Moses, Tucker, and Wu 2025) suggests that the source of evaluation may affect behavior differently, implying that responses under AI oversight need not mirror those under human monitoring. Avery, Leibbrandt, and Vecci (2023) provide field evidence that AI recruitment review changes behavior on the supply side by increasing women’s application completion rates. Complementing this, Almog, Lippman, and Martin (2026) show in an online real-effort experiment that workers change both the intensity and nature of their effort when they expect evaluation by AI rather than by a human. Iakovlev and Liang (2023) provide a theoretical rationale for such differences, emphasizing that human evaluators can tailor judgment to context, whereas algorithmic systems typically rely on a fixed set of covariates. Granulo, Fuchs, and Puntoni (2019) show that these distinctions may also operate through psychological channels: people exhibit different emotional and threat responses to replacement by robots or automation than to replacement by humans. More broadly, even apart from potential behavioral differences, AI expands the range of settings in which monitoring becomes operationally or economically feasible, since it can be cheaper, faster, or more effective than human oversight. We consider our setting to be one such case.

Lastly, we also contribute to the literature on attention and mistakes (Celsi and Olson 1988; Caplin 2023).¹⁰ We extend the canonical model of rational inattention with costs that scale linearly with Shannon mutual information (e.g., Caplin and Dean 2013; Matějka and McKay 2015). This extension allows us to incorporate two general features that we will later show fit well into our empirical setting. First, we allow for asymmetric costs of attention for different states. Second, we include behavioral factors in utility to accommodate the outcome of being overruled, which we will later refer to as the *AI oversight penalty*.¹¹ Bhattacharya and Howard (2022) find that the standard rational inattention model explains the equilibrium behavior of professional baseball players, and they estimate the linear cost of attention in this setting. Our rich data set, containing tournaments from both periods

¹⁰We acknowledge the ongoing debate regarding what to classify as a mistake (Nielsen and Rehbeck 2022). Here, we refer to a *mistake* as a decision that is incorrect ex-post relative to an objectively correct answer. It is fair to argue that tennis umpires are not making mistakes, as incorrect calls are just a result of the cost structure of becoming informed.

¹¹Other papers have examined how algorithmic recommendations directly influence decision-makers’ preferences. McLaughlin and Spiess (2024) propose a theoretical framework in which algorithmic recommendations not only shape beliefs but also establish a reference point. More specifically related to our work, Albright (2024) exploits a natural experiment in bail courts to isolate the causal effect of recommendations, independent of the prediction component, demonstrating that recommendations influence judges’ preferences by reducing the cost of making lenient decisions.

(with and without oversight), allows us to estimate not only the parameters associated with the cost of attention but also the state-dependent utility loss incurred when being overruled by AI. This novel element is a central component of our research question. Furthermore, our research introduces one of the first cognitive economic models that incorporate the influence of behavioral factors on rational attention allocation. Recent work highlights the value of expanding rational attention modeling to account for emotional states (Almog and Martin 2024; Bolte and C. Raymond 2023; Jerath and Ren 2021). We advance the attention literature by examining how AI tools affect attention allocation, a connection modeled by Boyacı, Canyakmaz, and de Véricourt (2023) and documented empirically by Grimon and Mills (2025) in the context of Child Protective Services’ assignment of investigations.

2 Setting and Data

In March 2006, at the Nasdaq-100 Open in Key Biscayne (currently known as the Miami Open), Hawk-Eye review was officially used for the first time at an ATP Tour event.¹² Later that year, many tennis tournaments that use non-clay surfaces, including the US Open, adopted this new technology, allowing players to challenge calls.¹³ Hawk-Eye uses six to ten computer-linked television cameras positioned around the court to collectively create a three-dimensional representation of the ball’s trajectory. Hawk-Eye performs with an average error of just 3.6 mm, so just like the ATP Tour we will consider Hawk-Eye readings to be the ground truth.¹⁴ While this error rate means that Hawk-Eye readings are not perfectly reflective of ground truth, our binning thresholds (20 mm, 100 mm) are far larger than this error rate, so our main qualitative patterns are unlikely to be driven by measurement noise in the ground truth.

We study the decision-making of tennis umpires before and after the introduction of Hawk-Eye review in professional tennis, focusing on the umpire’s judgment about whether a ball bounced in or out of bounds.¹⁵ This setting has several advantages for studying the impact of AI oversight on individual decision-making. First, the use of Hawk-Eye review in tennis was one of the pioneering uses of modern AI to conduct oversight in a work setting. Second, it is a setting of economic significance: the global revenue for the ATP Tour was

¹²The ATP Tour is the top tennis tour organized by the Association of Tennis Professionals.

¹³Players have a fixed number of challenges that they can make per match, but successful challenges do not count against this limit. Players were not allowed to challenge chair umpire decisions on non-clay courts before this system was put in place.

¹⁴For perspective, the standard diameter of a tennis ball is 67 mm.

¹⁵Our consolidated data set includes one tournament that was held without Hawk-Eye review after the 2006 Nasdaq-100 Open. However, this does not provide enough statistical power for running a difference-in-difference analysis. Hence, we use the shorthand that there is a *before* and an *after* period of Hawk-Eye review, even though this is strictly only true at the tournament level.

147.3 million USD in 2022, and the global revenue of Hawk-Eye was 71.6 million Euros in the 12 months to the end of March 2023.¹⁶ Third, there was a testing period during which the technology was used solely for broadcast and data recording purposes, without a challenge system in place, allowing us to track umpire performance both before and after the formal introduction of Hawk-Eye review in a tournament.¹⁷ Fourth, in this setting, there is an objectively correct decision and a reliable measure of it, which allows us to identify human mistakes. Finally, this setting offers the simplest possible decision-making problem, as the task is simply to match a binary action (call in or out) to a binary state (ball bounces in or out), which greatly simplifies our empirical and theoretical analysis.

2.1 Professional Tennis, Umpires, and Hawk-Eye

Professional tennis is a racket sport that is played either one-on-one (singles) or two-on-two (doubles). In singles tennis, two players compete against each other on opposite sides of the tennis court. The objective is to score points by hitting the ball “in” (within the bounds of the opponent’s side of the court). We will concentrate solely on men’s singles matches (singles matches where both players are men) because the analysis of other formats is underpowered in our data.¹⁸ The scoring system for tennis is based on a series of points, games, and sets. Players accumulate points to win a game and accumulate games to win a set. Matches are typically played as the best of three or five sets, with each set requiring a player to win at least six games.

A crew of up to ten umpires is involved in officiating a match. The roles typically include one chair umpire who oversees the entire match, making decisions on points, penalties, and overall match control. Additionally, there are usually nine line umpires positioned around the court, each responsible for specific lines, with the sole duty to determine whether the ball bounced in or out when it is relatively close to their respective line.¹⁹ The chair umpire has the authority to overrule the decisions of the line umpires if necessary, so drawing inferences on the performance of an individual umpire is complicated, especially since we do not have data on whether the chair umpire overruled an individual line umpire. Thus, in this paper, we evaluate the performance of the umpire crew as a whole. International chair umpires

¹⁶Sources: <https://www.breakingnews.ie/sport/revenues-at-hawk-eye-firm-increase-to-e71-6m-as-higher-costs-hit-profits-1534936.html> and <https://www.zippia.com/atp-tour-careers-985960/revenue/>.

¹⁷Our officiating source noted the following about the broadcast period: “*Even though these replays are not shown on the big screen, they inform the media and TV commentators, which can lead to widespread comments about the umpire being wrong.*” This suggests that we are likely underestimating the impact of AI oversight on umpire decision-making.

¹⁸If more data was available on the other formats, it might be of interest to examine whether the impacts of AI oversight vary by gender or team composition.

¹⁹Figure C.1 provides a visual representation of the positioning of each umpire.

are certified with Gold, Silver, or Bronze badges, while line umpires are graded according to national or other systems. Certification and grading methods remained unchanged during the study period, and Hawk-Eye metrics were not incorporated into these evaluations for the first couple of years. A leading tennis official who was involved in testing and introducing Hawk-Eye provided valuable institutional insights that will be used throughout the paper.

The Hawk-Eye review protocol works by endowing players with 2-3 challenges per set.²⁰ If a challenge is successful, players do not lose a challenge opportunity. When a player challenges a call, a computerized path and the final landing location of the ball are displayed on a large screen in the stadium for the umpire, players, and the crowd to observe the outcome of the challenge. The public nature of the challenge process adds excitement for the spectators, but simultaneously adds pressure to the umpires, as their decisions are being publicly scrutinized.

An important component of the review system implementation is the asymmetric resolution of incorrect calls, which varies depending on whether the challenged call was initially in or out. If a ball is initially ruled in by the umpires and the challenge successfully overturns the call, then the point ends with the challenger winning the point and the correct outcome being enforced. However, when a ball is initially ruled out by the umpires, but the review shows otherwise, enforcing the correct outcome is not always possible because the point was unnecessarily stopped. In this case, the umpire has to make an arbitrary decision on whether the opponent of the challenger had a real chance to return the ball. If the answer is yes, the point has to be replayed from scratch. This situation can be perceived as detrimental due to the inability to implement the correct outcome (continuing the point from where it stopped) and because it forces umpires to make arbitrary decisions that, in multiple instances, have been shown to infuriate one of the players involved.

2.2 Data

While this paper is not the first to utilize tennis data for drawing inferences on decision-making, we have assembled a novel and rich data set that, for the first time, enables us to comprehensively assess the performance of professional tennis umpires. To achieve this goal, we utilized three distinct data sources. Our primary data set, the *Hawk-Eye Base* data set, encompasses precise information on points and, crucially, the location of every bounce for over 100 tournaments.²¹ Our second data set, the *Challenge* data set, tracks the outcomes of challenges during a sample of tournaments that took place in the first few years following

²⁰In practice, players very rarely exhaust their challenges. This is advantageous for our research question, as umpires are almost always under the threat of being challenged.

²¹We excluded from the analysis the 15 clay court tournaments, we elaborate on this decision in the Appendix [A.1](#).

the implementation of Hawk-Eye review in professional tennis. In the period of time where challenges are used, the first data set only permits drawing conclusions on players’ behavior because, when observing a correct call, it is not possible to disentangle whether the umpire made the correct call initially or if a Hawk-Eye review corrected the umpire’s incorrect call. In contrast, the second data set allows precise identification of incorrect calls, as every won challenge is, by definition, an admission of the umpire’s mistake. Nonetheless, this second data set is incomplete because it only includes points that players decided to challenge, and this selection is susceptible to players’ perceptual and behavioral biases. Merging the first two data sets was challenging due to systematic inconsistencies in the Challenge data set (e.g., the score was often flipped across players and sometimes missing). To overcome this limitation, we turned to a third data source: a manual review of points by replaying match videos. This allowed us to identify the ground truth for challenges in a subset of matches, enabling us to determine an effective approach for merging the first two data sets.²² In the end, we were able to merge 2,038 out of the 2,108 challenges registered for those 28 tournaments (a 97% merge rate).

The consolidated data set produced by merging the Hawk-Eye Base data set and the Challenge data set comprises a total of 698 matches across 35 distinct tournaments, and summary statistics for this data set can be found in Table 1.²³ This data set includes 109 matches from seven tournaments that were played before Hawk-Eye review and 589 matches from 28 tournaments after Hawk-Eye review was active. The pre- and post-Hawk-Eye samples are very similar along the main point-level dimensions most relevant for identification, especially the shares of bounces within 20 mm and 100 mm of the line. The differences in tournament and round composition motivate the inclusion of match-level controls in future analyses.

We will now offer a more comprehensive description of the composition and role of each of the three data sources.

2.2.1 Hawk-Eye Base Data Set

The main source of data for this paper is the Hawk-Eye Base data set, which consists of the official Hawk-Eye data for matches played at the international professional level between March 2005 and March 2009. This data set provides a number of pieces of information for each point, including the position of every bounce captured by the Hawk-Eye system during that point, the serving player, the ongoing score, and the point winner. Altogether, this data

²²The video auditing process was also instrumental in validating the best rules for determining when a call was made incorrectly.

²³Table C.1 provides information on the month, category, court type, and the number of matches for each tournament.

	Before Hawk-Eye review (PostHK=0)	After Hawk-Eye review (PostHK=1)	Total	Post - Pre [Std. diff.]
<i>A. Sample counts</i>				
Players	71	161	174	
Tournaments	7	28	35	
Matches	109	589	698	
Points	15,439	83,898	99,337	
Bounces	71,638	402,446	474,102	
<i>B. Tournament tier (share of tournaments)</i>				
ATP 250	28.6%	42.9%	40.0%	+14.3 pp [0.30]
ATP 500	0.0%	10.7%	8.6%	+10.7 pp [0.49]
ATP 1000	71.4%	46.4%	51.4%	-25.0 pp [0.53]
<i>C. Match round (share of matches)</i>				
Final	6.4%	4.6%	4.9%	-1.8 pp [0.08]
Semifinal	12.8%	9.0%	9.6%	-3.8 pp [0.12]
Quarterfinal	17.4%	14.9%	15.3%	-2.5 pp [0.07]
Round of 16	12.8%	19.2%	18.2%	+6.3 pp [0.17]
Other	50.5%	52.3%	52.0%	+1.8 pp [0.04]
<i>D. Bounce composition (share of all bounces)</i>				
< 20 mm from the line	0.78%	0.69%	0.70%	-0.09 pp [0.01]
< 100 mm from the line	3.66%	3.53%	3.55%	-0.13 pp [0.01]
Serves	29.2%	27.6%	27.8%	-1.64 pp [0.04]
Non-serves	70.8%	72.4%	72.2%	+1.64 pp [0.04]
<i>E. Average speed in km/h (s.d.)</i>				
Serves	147.7 (24.2)	150.0 (23.1)	149.6 (23.3)	+2.3 [0.10]
Non-serves	81.7 (21.4)	83.3 (20.4)	83.1 (20.5)	+1.6 [0.08]
All	101.0 (37.4)	101.7 (36.5)	101.6 (36.6)	+0.7 [0.02]

Table 1: Summary statistics for the consolidated data set. The last column reports the difference in the sample characteristics before and after Hawk-Eye review. Bracketed values are absolute standardized differences.

set includes information on 1.8 million bounces from more than 1,800 men’s singles matches, spanning various prestigious tournaments such as Grand Slam tournaments, Masters 1000 tournaments, and International series tournaments. This data set has been used in the past to study important economics questions such as risk management (Ely, Gauriot, and Page 2017), tournament incentives (Gauriot and Page 2019), and mixed strategy play (Gauriot, Page, and Wooders 2023). These papers have focused on the players’ perspective, and the Hawk-Eye Base data set is well-suited for this purpose. However, this data set is insufficient to analyze umpires’ performance, as it does not permit us to identify whether the final call comes from the umpire or via an overturned challenge.

2.2.2 Challenge Data Set

The Challenge data set was originally obtained from the ATP Tour and encompasses all challenges recorded during 35 tournaments across the initial three years following the introduction of Hawk-Eye (2006-2008). It captures comprehensive information on 2,784 challenges, including details about the players involved, the match itself, the specific point in the match when the challenge was made, and the outcome of each challenge. Of the 35 tournaments, we use the 28 that also appear in our Hawk-Eye Base dataset. In order to compile this data set, Abramitzky, Einav, Kolkowitz, and Mill (2012) undertook the arduous effort of collecting and compiling umpire match sheets. Abramitzky, Einav, Kolkowitz, and Mill (2012) acknowledge the selection problem involved in only observing the challenged points (around 2.6% of total points). Nevertheless, they are able to draw inferences on the optimality of decision-making from the players’ standpoint by relying on the idea that a higher propensity to challenge implies challenging less conservatively. However, without the point-by-point data, it is hard to assess the umpire’s performance once Hawk-Eye review was active. By merging the first two data sets, we solve the aforementioned selection problem and gain information on what was originally called by the umpire crew.

2.2.3 Video Auditing

In order to assess the quality of our merging algorithm, we audited full-match video replays using TennisTV, the official ATP streaming service. We identified 43 matches that appeared in both the Hawk-Eye Base and Challenge data sets, providing us with an assessment of how well a given merging algorithm would work in the 144 challenges witnessed during these matches.

Using variables such as match, point, distance, and who hit the ball, we developed a merging algorithm with a 99.3% merge rate, for which only 4.9% of the video-audited challenges were merged to an incorrect point. The algorithm consists of eight iterations,

starting with the most stringent merging rule and gradually relaxing criteria for challenges that persist without merging.²⁴

As an additional benefit, auditing match replays enabled us to test the effectiveness of the four criteria we implemented to identify incorrect calls (two criteria for each type of mistake). We deem an in call to be incorrect when a player’s stroke is recorded as bouncing out in the Hawk-Eye Base data set, yet they still win the point, or if the data set records at least three more strokes after an out bounce. We deem an out call to be incorrect if all the strokes of a point are recorded as in, and the player delivering the final hit loses, or if there is a second serve after the first one was recorded as in the Hawk-Eye Base data set.²⁵

3 Empirical Results

In this section, we use our consolidated data set to study the impact of AI oversight on umpiring decision-making in professional tennis.²⁶ We begin by examining changes in the overall mistake rate for all hits. We then highlight the influence of distance from the line and whether a ball was in or out.

3.1 Overall Mistake Rate

Before Hawk-Eye review was introduced, the umpire mistake rate was only 0.61% of all ball bounces. However, if we look at the 2,622 (3.6%) bounces that fell within 100 mm of the line, the mistake rate jumps to 13.89%. Looking at the 556 (0.78%) bounces that fell within 20 mm of the line, the mistake rate jumps even more to 32.91%. Given our interest in the impact of AI oversight on human mistakes, our primary focus will be on calls within 100 mm or 20 mm of the line.²⁷

While the likelihood of a close call (a ball bouncing within 100 mm or 20 mm of the line) did not change substantially after the introduction of Hawk-Eye and the probability of a close call being in or out did not change substantially either,²⁸ we find that the mistake rate

²⁴A detailed explanation of the merging algorithm is provided in Appendix A.2.

²⁵Appendix A.3 documents these criteria further.

²⁶Our officiating expert confirmed that the pool of umpires remained stable throughout the study period. This eliminates the possibility that the estimated effects are due to changes in the umpire sample. Additionally, suspensions and terminations are implemented solely for disciplinary reasons rather than performance, which further alleviates this concern.

²⁷We also find that over 93% of the challenges in our sample are for calls within 100 mm of the line.

²⁸See Table 1 for the likelihood of a close call. The probability that a ball was out when it bounced within 100 mm of the line was 45.3% before Hawk-Eye review and 46.8% after, and the corresponding numbers for 20 mm are 48.4% and 48.9%.

on close calls did change substantially. We estimate the effect of AI oversight on umpires’ performance using the following specification:

$$\mathbb{1}(\textit{Incorrect Call})_{ipm} = \alpha_0 + \alpha_1 \textit{PostHK}_m + \beta X_p + \gamma Y_m + \epsilon_{ipm} \quad (1)$$

$\mathbb{1}(\textit{Incorrect Call})$ is an indicator variable equal to 1 if the call i made by the umpire in point p in match m was incorrect. \textit{PostHK} is an indicator variable equal to 1 if the match m belongs to a tournament with the Hawk-Eye review active. X is a vector of point characteristics: distance fixed effects (by 20 mm bins), speed (whether it is below or above the median), score, game, set and an indicator if the point played is in the tie-break stage.²⁹ Y is a vector of match characteristics that includes round and tournament tier.³⁰

Our main coefficient of interest is α_1 , which can be interpreted with OLS regression as the effect in percentage points of AI oversight on the probability that an umpire’s call is incorrect. Table 2 reports the results of estimating this equation for calls within 100 mm of the line. In our primary specification, the mistake rate is estimated to decrease by 8% (1.1 p.p.). This decrease in the mistake rate is consistent with a model of rational inattention in which the psychological costs of being overruled by the AI outweigh any benefits, such as attentional free-riding (see Section 4 for a model of this form).

Importantly, the average effect of Hawk-Eye review is fairly stable across specifications. Adding point and match controls changes the estimate only from -1.4 to -1.1 percentage points, so the main change with added controls is a loss of precision rather than a meaningful loss in magnitude. The noise in these estimates is unsurprising because the single \textit{PostHK} coefficient averages over sharply heterogeneous effects by distance to the line and by which side of the line the ball lands on. First, an important component to call difficulty is the distance from the line, so the ability to change attention enough to fix a mistake might depend on the distance from the line. Second, it might matter which side of the line the ball is on, and umpires might have different perceptions of the cost of Type I and Type II errors after the introduction of Hawk-Eye review. As subsequent figures show, review lowers mistake rates for many close calls but increases them for the narrow region just outside the line, so the overall average masks offsetting effects.

3.2 Shift in Type I and Type II Errors

Figure 1 illustrates the relationship between mistake rates and the distance to the line for the 16,812 bounces within 100 mm of the line. In this figure, each dot represents the rate at

²⁹A tie-break is a one-off game held to decide the winner of a set when two players are locked at 6-6.

³⁰Given that the period without oversight has no 500-tier tournaments, we aggregate the 250 and 500 groups, so we basically control for whether the match is a Masters 1000 tournament or not.

	(1)	(2)	(3)	(4)
	Incorrect call			
PostHK	-0.014** (0.007)	-0.011 (0.007)	-0.011 (0.007)	-0.011 (0.007)
Point controls		X	X	X
Match controls			X	X
Cluster level				Match
N	16,812	16,812	16,335	16,335
Baseline mean	.139	.139	.136	.136

Standard errors in parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 2: OLS regressions of umpire mistakes for balls bouncing within 100 mm of the line. $PostHK = 1$ if the Hawk-Eye review is active; point controls include distance fixed effects (by 20 mm bins), speed (whether it is below or above the median), score, game, set, and an indicator if the point played is in the tie-break stage; and match controls include round and tournament tier.

which umpires made an incorrect call for all balls bouncing within one of ten 20 mm bins. The five bins to the left of the dashed line correspond to balls that landed out of bounds, while the five bins to the right side of that line represent balls that landed inside of the line. For the rest of the paper, negative distances will be used to denote the distance to the line for balls that bounced outside.

The blue dots in the figure were calculated using tournaments before the introduction of Hawk-Eye review. The red dots represent tournaments with Hawk-Eye review. While it is expected that umpires' performance decreases as the ball gets closer to the line, we can also observe that the red line lies mostly under the blue line (this holds true for eight of the ten bins). This mirrors the results in Table 2, which shows that the overall performance of umpires improves after the introduction of AI oversight, and in Table 6a, which shows that this is particularly true for balls bouncing between 100 and 20 mm away from the line. In addition, it is noticeable that the only region of the court where the introduction of Hawk-Eye increased the rate at which umpires make mistakes is the two closest bins on the outside of the line.

To further analyze this change in the type of errors made, we examine how the impact of AI oversight on umpire performance varies with distance to the line. We re-estimate equation 1 (including point and match controls), but now we interact $PostHK$ with the indicator variables for each of the ten 20 mm distance bins into which the ball bounced. We report the estimated coefficients of these interaction terms graphically in Figure 2. For each bin, we plot the estimated coefficient as a dot and provide the respective 95% confi-




Figure 1: Incorrect call rates by proximity to the line. Each dot is the rate of incorrect calls for a bin of 20 mm. Dots to the left of the dashed line represent bins out of bounds, and those to the right of the dashed line represent bins in bounds.

dence interval around the estimated coefficient. The red dashed line, which separates the in and out parts of the court, shows a sharp discontinuity. The first bin to the left of the red dashed line (balls bouncing just out) exhibits the most significant positive increase in incorrect calls, with an estimated coefficient of 8.6 percentage points (significant at the 1% level). From that point onward, the coefficients gradually adjust back down to the average treatment effect, just below zero. In the first two bins to the right of the red dashed line (balls bouncing in), we observe the most substantial decrease in incorrect calls, with estimated coefficients of -4.1 (significant at the 5% level) and -7.7 (significant at the 1% level) percentage points. Similarly, as we continue to move to the next bins in that direction, the estimated coefficients gradually increase to the average treatment effect level. As highlighted by Goldsmith-Pinkham, Hull, and Kolesár (2024), our estimates from Figure 2 may suffer from “contamination bias” due to the regression of multiple treatments while controlling for observables in an additively separable manner. To address this concern, we implement one-treatment-at-a-time regressions, one of the suggested solutions by Goldsmith-Pinkham, Hull, and Kolesár (2024). Figure C.2 shows that our original estimates do not suffer from this type of bias.

We provide an event-study version of Figure 2 in Figure 3. This figure illustrates how the impact of AI oversight on umpire performance progresses over time by separately examining




Figure 2: Impact of AI oversight on the incorrect call rate by proximity to the line. Each dot represents the coefficient on the interaction between the distance bin and PostHK, the indicator variable that equals 1 if the Hawk-Eye review system is active. Controls for point and match characteristics are included.




Figure 3: Impact of AI oversight on the incorrect call rate by proximity to the line (by time sub-period). Matches are grouped into those with Hawk-Eye review in all of 2006 and the first semester of 2007 and those with Hawk-Eye review in the second semester of 2007 and all of 2008. Each dot represents the coefficient on the interaction between the distance bin and PostHK, the indicator variable that equals 1 if the Hawk-Eye review system is active.

matches in the first half and second half of the period we study. The first group comprises tournaments from 2006 and the first semester of 2007, while the second group includes tournaments from the second semester of 2007 and those from 2008. Figure 3 suggests that the shift in umpires’ error types was not immediate, but instead took over a year to fully manifest. We chose this division into two time sub-periods to yield a similar number of observations in both groups, and the results are robust to breaking the tournaments into more time sub-periods, though the effects become more noisily estimated.

Clearly, the key shift in errors occurs for the closest calls. What drives this shift? Figure 4 shows the tournament-specific rates at which umpires called a ball in when it bounced within 20 mm of the line, regardless of which side it bounced. Each of the 28 dots represents the rate at which balls were called in for a particular tournament that allowed AI review, sorted by time. The first observation is that this rate increased from 42.8% to 49% with the introduction of AI review. In 22 out of the 28 tournaments with AI review, umpires called the ball in more frequently than the mean of 42.8% observed in tournaments without AI review.³¹ Figure 4 suggests a gradual increase in the frequency of inside calls by umpires

³¹The fact that a majority of bounces close to the line are called out is notable given that it runs against

3.3 Alternative Mechanisms

Several alternative mechanisms could, in principle, contribute to the behavioral responses we document following the introduction of Hawk-Eye review. One possibility is institutional guidance or supervisory pressure to adopt a particular default for close calls. However, the leading official we contacted, who played a key role in the deployment of Hawk-Eye in the ATP, confirmed that umpires were not explicitly instructed to call more ‘ins’ for close calls. A second possibility is strategic adaptation driven by asymmetric challenge exposure. In our setting, challenge rates are similar across error types (Type I: 41.5%; Type II: 44.9%), suggesting that differential challenge likelihood is unlikely to be the primary driver of the shift in the types of errors we observe.³³ Finally, one might worry that the effects reflect reputational concerns driven by increased information about umpire performance. Yet even prior to the stadium-based review regime, incorrect calls were already visible to outside observers through broadcast coverage; the key change with Hawk-Eye is the public, on-the-spot overturning of on-court decisions. Taken together, these considerations point away from explanations based solely on directives, information flow, or sharply asymmetric review exposure, and instead motivate our modeling approach: in Section 4 we capture the introduction of AI oversight as creating an additional psychological cost associated with being overturned under review, which provides a coherent explanation of the observed reduction in overall mistakes alongside a systematic shift in error composition.

4 Recovering the AI Oversight Penalty

Finally, we structurally estimate the psychological costs of being overruled by AI using a model of rational inattentive umpires.³⁴ Here, we apply the rational inattention theory introduced by Sims (2003) and characterized by Matějka and McKay (2015) and Caplin and Dean (2013), which we refer to as the “Shannon model.” Bhattacharya and Howard (2022) use this model to study attention constraints in baseball, while Brown and Jeon (2024) use it to study consumers choosing among complex health insurance plans.

In our model, the set of states is $\omega \in \{\omega^I, \omega^O\}$ for the ball being in (ω^I) or out (ω^O), and the set of actions is $a \in \{a^I, a^O\}$ for calling the ball in (a^I) or out (a^O). The probability an incorrect call is challenged is η^I when the ball is in, and η^O when the ball is out, and both are equal to zero in tournaments without AI review. When the umpire’s call is not challenged, we assume that they receive a normalized utility of 0 when correct and -1 when incorrect.

³³Nevertheless, in Section 4 we allow challenge probabilities to differ by error type in the structural estimation.

³⁴While it could be interesting to consider the dynamic effects of being overruled by AI, we start by considering a static model.

When the umpire’s call is challenged, we also assume that they receive a normalized utility of 0 when correct (when their call is upheld). But when incorrect (their call is overturned), we assume that they receive c^I when the ball is in and c^O when it is out. Thus, before taking into account attentional costs, gross expected utility U is

$$U(a, \omega) = \begin{pmatrix} \omega_I & \omega_O \\ 0 & -1 + \eta^O (1 + c^O) \\ -1 + \eta^I (1 + c^I) & 0 \end{pmatrix} \begin{matrix} a_I \\ a_O \end{matrix}$$

We parameterize the utility of an overturned call in this way so that we can interpret c^I and c^O as the disutility of having been caught making a mistake by the AI system. For shorthand, we refer to these parameters as the *AI oversight penalty* when the ball is in or out. We expect that c^I and c^O are less than or equal to 0 because for the parameters to be positive the umpire would need to be happier having their incorrect call overturned than being correct in the first place. When c^I and c^O are equal to zero (when there is no AI oversight penalty), then the umpire receives a utility of 0 after their call is overturned. This is as if they only care that the correct call is implemented (due to altruism, relief, etc.) and do not care at all about having been caught making a mistake by the AI system (due to shame, embarrassment, subsequent arguments, etc.). When c^I and c^O are equal to -1 , then the umpire receives a utility of -1 after their call is overturned. This is as if any utility gain from having the correct call implemented is perfectly offset by the AI oversight penalty. When c^I and c^O are less than -1 , the AI oversight penalty dominates any utility gain from having the correct call implemented.

As is standard in models of attention, we assume that the umpire starts out with prior belief μ , where $\mu(\omega)$ is the probability of state $\omega \in \{\omega^I, \omega^O\}$. After receiving a noisy mental signal of whether the ball is in or out, the umpire forms posterior $\gamma \in \Gamma$, where $\gamma(\omega)$ is the probability of state $\omega \in \{\omega^I, \omega^O\}$. Given a posterior belief γ , the umpire decides what call to make (mixing between actions is allowed), and we assume the umpire maximizes expected utility when making a choice of how often to take each action at that posterior.

Under the assumption that the umpire updates their prior belief using Bayes rule, their attention can be represented by a Bayes-consistent information structure π , which stochastically generates posterior beliefs. Specifically, π is a function that maps a state ω from a set of states Ω into $\Delta(\Gamma)$, the set of probability distributions over Γ that have finite support, so that $\pi : \Omega \rightarrow \Delta(\Gamma)$. Let Π denote the set of all such functions, $\pi(\gamma)$ be the unconditional probability of posterior $\gamma \in \Gamma$, $\pi(\gamma|\omega)$ be the probability of posterior γ given state ω , and $\Gamma(\pi) \subset \Gamma$ denote the support of a given π . We limit the set of information structures to

those in $\Pi(\mu) \subset \Pi$ that generate correct posteriors for a given prior belief μ , so that

$$\Pi(\mu) = \left\{ \pi \in \Pi \mid \forall \gamma \in \Gamma(\pi), \forall \omega \in \Omega, \gamma(\omega) = \frac{\mu(\omega)\pi(\gamma|\omega)}{\sum_{\omega \in \Omega} \mu(\omega)\pi(\gamma|\omega)} \right\} \quad (2)$$

Next, we assume information structures are chosen, which we interpret as the umpire choosing how much attention to pay and what aspects of the game to pay attention to, as in Bhattacharya and Howard (2022). Each information structure $\pi \in \Pi(\mu)$ has an additively-separable cost $K(\pi)$ that scales linearly with the Shannon mutual information between posteriors and the prior. Formally, K is determined by the function

$$K(\pi, \kappa, \mu) = \kappa \left(\left[\sum_{\gamma \in \Gamma(\pi)} \pi(\gamma) \sum_{\omega \in \Omega} [\gamma(\omega) \ln(\gamma(\omega))] \right] - \sum_{\omega \in \Omega} [\mu(\omega) \ln(\mu(\omega))] \right) \quad (3)$$

where $\kappa \in \mathbb{R}_{++}$ is a linear cost parameter, which can be interpreted as the marginal cost of attention. We assume that κ is increasing in the difficulty of judging where a ball landed and assume that the difficulty of this perceptual task is unimpacted by the introduction of Hawk-Eye review.

We extend the standard Shannon model by allowing for different costs of attention for different states, so that

$$K(\pi, \kappa, \mu) = \kappa^I \left(\left[\sum_{\gamma \in \Gamma(\pi)} \pi(\gamma) \gamma(\omega^I) \ln(\gamma(\omega^I)) \right] - \mu(\omega^I) \ln(\mu(\omega^I)) \right) \quad (4)$$

$$+ \kappa^O \left(\left[\sum_{\gamma \in \Gamma(\pi)} \pi(\gamma) \gamma(\omega^O) \ln(\gamma(\omega^O)) \right] - \mu(\omega^O) \ln(\mu(\omega^O)) \right) \quad (5)$$

where $\kappa^I, \kappa^O \in \mathbb{R}_{++}$. Whitney, Wurnitsch, Hontiveros, and Louie (2008) use data from Wimbledon 2007 to show that umpires call more tennis balls as being out (when actually in) than in (when actually out). Their results are consistent with psychometric experiments that document how moving objects generate a perceptual bias and are perceived as being shifted in the direction of their motion. Extending the Shannon model by allowing different costs of attention for different states would help to accommodate these types of biases more broadly.

With the Shannon model, the optimal information structure has one posterior at which each action is taken. We denote these posteriors as γ^I when call in and γ^O when call out. By the Invariant Likelihood Ratio (ILR) of Caplin and Dean (2013), these optimal posteriors must obey

$$\frac{\gamma^I(\omega^I)}{\gamma^O(\omega^I)} = e^{\frac{U(a^I, \omega^I) - U(a^O, \omega^I)}{\kappa^I}} \quad (6)$$

$$\frac{\gamma^O(\omega^O)}{\gamma^I(\omega^O)} = \frac{1 - \gamma^O(\omega^I)}{1 - \gamma^I(\omega^I)} = e^{\frac{U(a^O, \omega^O) - U(a^I, \omega^O)}{\kappa^O}} \quad (7)$$

For matches where there was no AI review, the ILR condition, given by (6) and (7), allows us to express the marginal costs of attention κ^I and κ^O cleanly as

$$\kappa^I = \frac{1}{\ln \gamma^I(\omega^I) - \ln \gamma^O(\omega^I)} \quad (8)$$

$$\kappa^O = \frac{1}{\ln \gamma^O(\omega^O) - \ln \gamma^I(\omega^O)} \quad (9)$$

When there is AI review, (6) and (7) allow us to solve for the AI oversight penalty, which is the umpire’s disutility from having been caught making a mistake by the AI system. Those values are given by

$$c^I = \frac{1 - \kappa^I (\ln \gamma^I(\omega^I) - \ln \gamma^O(\omega^I))}{\eta^I} - 1 \quad (10)$$

$$c^O = \frac{1 - \kappa^O (\ln \gamma^O(\omega^O) - \ln \gamma^I(\omega^O))}{\eta^O} - 1 \quad (11)$$

4.1 Structural Estimates

To structurally estimate the AI oversight penalty, we undertake two steps. First, we estimate the marginal costs of attention when there is no AI review. When there is no AI review, these marginal costs are fully determined by the optimal posteriors $\gamma^I(\omega^I)$ and $\gamma^O(\omega^O)$. Because there is a single posterior for each action in our model, the optimal posteriors are equal to the *revealed posterior*, which is $P(\omega|a)$, the probability of each state conditional on the action taken. Our estimate of the *true* revealed posterior is the *observed* probability of each state conditional on the action taken. We present the estimates recovered from the data, along with the bootstrap standard errors obtained from drawing 1000 random samples with replacement, each containing the same number of observations as the original dataset.

As shown in Table 3, our estimate of the revealed posterior for the ball being in when calling in is 0.849, and our estimate of the revealed posterior for the ball being out when calling out is 0.876 for all calls within 100 mm of the line before the introduction of AI oversight. Based on (8) and (9), the marginal costs of attention that rationalize these values are $\kappa^I = 0.580$ and $\kappa^O = 0.510$.

As a second step, we use these estimated costs of attention and the observed challenge rates (our estimates of the true challenge rates η^I and η^O) to estimate the AI oversight penalties. As shown in Table 4, the observed challenge rates are 0.449 when balls were in and 0.415 when balls were out. Given the values of κ^I and κ^O determined without AI review,

Parameter	Recovery equation	Estimate
$\gamma^I(\omega^I)$	N/A	0.849
$\gamma^O(\omega^O)$	N/A	0.876
κ^I	8	0.580 (0.025)
κ^O	9	0.510 (0.024)

Standard errors in parentheses

Table 3: Optimal posteriors and estimated costs of attention before the introduction of AI oversight.

the rationalizing AI oversight penalties are $c^I = -1.374$ and $c^O = -0.903$. It is important to highlight that the AI penalty, in our case, can be seen as a lower bound on the costs of shaming. Under the scenario that umpires do not care at all about the final outcome being correctly implemented, then the cost of shaming equals the AI oversight penalty, but otherwise, the shaming effect is bigger. These results are consistent with anecdotal evidence regarding the increased controversy surrounding this type of error.

Parameter	Recovery equation	Estimate
η^I	N/A	0.449
η^O	N/A	0.415
c^I	10	-1.374 (0.121)
c^O	11	-0.903 (0.116)

Standard errors in parentheses

Table 4: Challenge rates and estimated AI oversight penalties.

These estimates suggest that the introduction of AI oversight left the utility of Type I errors on overturned calls largely unchanged, as there is an increase from -1 without Hawk-Eye review to -0.903 after, but the difference is not statistically significant. On the other hand, the utility of Type II errors on overturned calls reduces sharply, from -1 without Hawk-Eye review to -1.374 after. Here, the AI oversight penalty dominates any utility gain from having the correct call implemented. Thus, this suggests that with AI oversight umpires care approximately 37% more about Type II errors on overturned calls.

5 Examining Heterogeneity

In this section, we study three main sources of heterogeneity: skill, stakes, and type of perceptual task. Practically, this relates to differences by tournament stage (finals or earlier stage), tournament tier (higher or lower tier), and shot type (serve and non-serve).

5.1 Heterogeneous Response by Umpire Skill

Umpires are rewarded for good performance with assignments to advanced-stage matches, as the pool of officials needed becomes smaller and organizers can be more selective as the tournament progresses. Our officiating source confirmed that not all umpires (particularly line umpires) are professional officials, and you can expect to see more of them in the early stages. These selection dynamics within tournaments create variation in umpiring skill. Using later tournament stages to identify more skilled umpires may confound the analysis by including matches with higher stakes. However, we demonstrate that this is not the case at the end of the section.

To examine whether the impact of AI oversight differs by tournament stage, we split our sample into two groups: the first group includes Final, Semifinal, and Quarterfinal matches, where we expect to find the most skilled umpires, while the second group contains the remaining matches. In Figure 5, we examine the AI oversight effect separately on the two groups of matches we described. The change in the type of errors in bins closer to the line is more pronounced and only significant for matches in the early stages, as highlighted by the magenta markers. The estimated effects for advanced-stage matches (green markers) are not statistically significant. Despite having fewer observations, the estimated coefficients are also smaller in absolute terms in bins around the dashed line.

In a study on radiologists, Chan, Gentzkow, and Yu (2022) found that aversion to false negatives tends to be negatively related to radiologist skill. Specifically, lower-skilled radiologists are more likely to falsely diagnose a healthy patient with pneumonia (false positive) because this error is less costly than missing a diagnosis. Our findings align with those of Chan, Gentzkow, and Yu (2022), as less-skilled umpires are more likely to avoid the more costly type of error (calling a ball out when it is actually in).

It is important to note that advanced matches also involve higher stakes. It is challenging to disentangle whether the observed effects are due to lower-skilled umpires or higher stakes. To address this, we can compare differences in stakes by examining highly regarded tournaments, such as Masters 1000 events, against lower-tier tournaments like the 500 and 250 series. Figure C.5 shows that the discrepant patterns we observed between early and advanced-stage matches do not hold when we compare higher and lower-ranked tournaments.




Figure 5: Impact of AI oversight on the incorrect call rate by proximity to the line (by tournament stage). Matches are grouped into those at the final, semifinal, and quarterfinal stages and those at all other (earlier) stages of a tournament and then analyzed separately. Each dot represents the coefficient on the interaction between the distance bin and PostHK, the indicator variable that equals 1 if the Hawk-Eye review system is active.

If the difference in stakes is what drives the heterogeneity observed in Figure 5, we should expect to see the same patterns between lower and higher-ranked tournaments, but we do not. A natural concern is that higher-tier tournaments employ better officials. As our officiating expert noted, assignment decisions reflect many factors beyond tier (e.g., geography, budget, and logistical constraints), and there is substantial overlap in the officiating pool across tiers. We therefore use tournament tier to proxy for match stakes, and we show that our conclusions are robust to specifications that account for observable match characteristics and persistent differences across officials.

5.2 Serves and Non-Serves

There is value in distinguishing between serves and non-serves,³⁵ as we can think of them as distinct tasks for the umpires. The primary differences, from an officiating standpoint, between serves and subsequent hits lie in speed and anticipation. Serves may be considered more challenging due to their significantly higher speed: the average speed for a serve in our sample is 156 km/h, compared to 82 km/h for non-serves.³⁶ However, serves have the advantage of anticipation: umpires know the ball is about to be served and will most likely bounce in a very specific region of the court (the serving box). This appears to be important: when comparing the slowest quartile of serves (averaging 133 km/h) with the fastest quartile of non-serves (averaging 106 km/h), we find that the error rates for serves are lower, despite still containing faster strokes. The average incorrect rate during the period without Hawk-Eye review for the first quartile in serves and the fourth quartile in non-serves is 11.6% versus 15.6% for balls within 100 mm of the line, and 27.5% versus 37.1% within 20 mm of the line. Therefore, even though serves are much faster than non-serves, there is compelling evidence to support that the anticipation advantage benefits umpires during serves.

In Figure 6, we break down the incorrect call rates by serve and non-serves. This figure shows that the shift in types of errors is present in both serves and non-serves. For serves, the dominant effect is the shift in errors. For non-serves, the dominant effect is a decrease in mistakes. This suggests that non-serves may have a higher potential for wholesale improvement, which could be due to their more manageable speed and the opportunity they provide to rectify mistakes caused by distraction or sparse attention.

Table 5 reports the results from estimating equation 1 separately for serves and non-serves that bounced within 100 mm of the line. In our main specification, we do not find evidence that AI oversight impacted umpire performance during serves. However, we do find evidence that AI oversight corresponds to a 2.3 percentage point decrease in incorrect calls for non-serves. This corresponds to a 17% reduction compared to the baseline mean level of

³⁵The serve starts off every point in tennis, with players alternating serving each game.

³⁶Figure C.6 shows the speed distribution for both types.

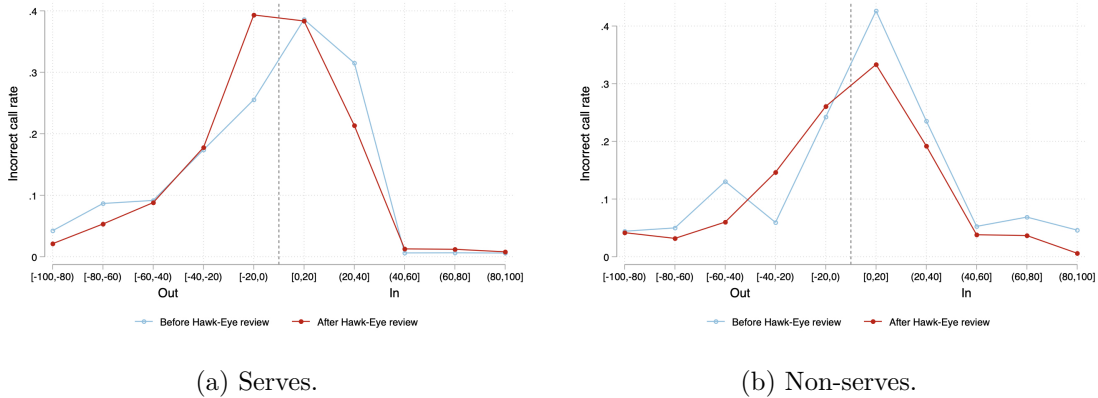


Figure 6: Incorrect call rates by proximity to the line (separately for serves and non-serves). Each dot is the rate of incorrect calls for a bin of 20 mm, and dots to the left of the dashed line represent bins out of bounds, and those to the right of the dashed line represent bins in bounds.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Incorrect call (Serves)				Incorrect call (Non-serves)			
PostHK	-0.006 (0.010)	-0.002 (0.009)	0.000 (0.009)	0.000 (0.009)	-0.024** (0.010)	-0.021** (0.010)	-0.023** (0.010)	-0.023** (0.011)
Point controls		X	X	X		X	X	X
Match controls			X	X			X	X
Cluster level				Match				Match
N	9,253	9,253	8,990	8,990	7,559	7,559	7,345	7,345
Baseline mean	.139	.139	.137	.137	.138	.138	.136	.136

Standard errors in parentheses
 * $p < .10$, ** $p < .05$, *** $p < .01$

Table 5: OLS regressions of umpire mistakes for balls bouncing within 100 mm of the line separately for serves and non-serves. $PostHK = 1$ if the Hawk-Eye review is active; point controls include distance fixed effects (by 20 mm bins), speed (whether it is below or above the median), score, game, set, and an indicator if the point played is in the tie-break stage; and match controls include round and tournament tier.

0.136.

If we just consider those calls when the ball bounces between 20 and 100 mm from the line, Table 6a shows that there is a statistically significant reduction in mistakes for both serves and non-serves. However, we find something different if we look at the closest calls. In Table 6b, we present the estimates for equation 1 when restricting the sample to bounces within 20 mm. This subset of calls is arguably the most complicated for umpires, and improving performance further would entail excessive cognitive costs. For these calls, the introduction of Hawk-Eye resulted in a 7.3 percentage point increase in incorrect calls for serves, representing a 22.9% increment from the baseline. While this result focuses on a very specific subset of the sample (serves within 20 mm), it provides evidence of which type of situations can lead AI oversight to backfire.

6 Conclusion

This paper studies how workers respond when an algorithm can publicly overturn their decisions. Using the introduction of Hawk-Eye review for close line calls in professional tennis, we document that AI oversight changes not only average accuracy, but also the composition of mistakes. On average, oversight improves performance on close calls: for bounces within 100 mm of the line, the incorrect-call rate falls by about 1.1 percentage points (roughly 8% relative to baseline). Yet the response is not uniform across states. For the closest decisions, umpires shift toward calling balls in. This threshold shift reduces Type II errors (calling out when in) but increases Type I errors (calling in when out), generating a clear backfire region: for balls just outside the line, mistake rates rise sharply even as they fall for balls just inside.

We interpret these patterns through a rational inattention model augmented with a utility loss from being overturned: the “AI oversight penalty”. Structural estimates imply that being overturned is more costly when the ball is in than when it is out, consistent with the institutional asymmetry that Type II errors can force a discretionary remedy, and that umpires care approximately 37% more about Type II errors on decisions that were overturned.

For organizations deploying AI oversight (audit-and-override systems, not just recommendations), three implications follow. First, accuracy metrics that collapse across error types can be misleading. Oversight may reduce overall error while shifting errors toward states where the organization or customers bear higher costs. Effective governance therefore requires monitoring error *composition* and cost-weighted performance, not only average accuracy. Second, the design of the override protocol is central. When correction is more visible, more disruptive, or more contentious for one error type, workers will rationally adjust

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Incorrect call (Serves)				Incorrect call (Non-serves)			
PostHK	-0.020** (0.008)	-0.018** (0.008)	-0.017** (0.008)	-0.017** (0.009)	-0.012 (0.009)	-0.017** (0.009)	-0.020*** (0.009)	-0.020** (0.011)
Point controls		X	X	X		X	X	X
Match controls			X	X			X	X
Cluster level				Match				Match
N	7,499	7,499	7,238	7,238	6,047	6,047	5,785	5,785
Baseline mean	.092	.092	.091	.091	.082	.082	.080	.080

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

(a) For balls bouncing 20-100 mm away from the line.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Incorrect call (Serves)				Incorrect call (Non-serves)			
PostHK	0.063** (0.031)	0.063** (0.031)	0.073** (0.032)	0.073** (0.030)	-0.035 (0.031)	-0.037 (0.031)	-0.031 (0.033)	-0.031 (0.034)
Point controls		X	X	X		X	X	X
Match controls			X	X			X	X
Cluster level				Match				Match
N	1,804	1,804	1,752	1,752	1,512	1,512	1,470	1,470
Baseline mean	.325	.325	.319	.319	.333	.333	.325	.325

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

(b) For balls bouncing within 20 mm of the line.

Table 6: OLS regressions of umpire mistakes (separately for serves and non-serves). $PostHK = 1$ if the Hawk-Eye review system is active; point controls include distance fixed effects (by 20 mm bins), speed (whether it is below or above the median), score, game, set, and an indicator if the point played is in the tie-break stage; and match controls include round and tournament tier.

thresholds toward the “safer” mistake. Designing symmetric remedies and reducing discretionary, high-friction corrections can mitigate perverse substitution. Third, oversight is most likely to distort behavior when tasks operate near human perceptual limits or when further accuracy improvements are cognitively expensive. In such settings, complementary changes (training, workflow redesign, better escalation rules, or selectively automating edge cases) may dominate simply adding review.

More broadly, the tennis setting illustrates a practical reality: even near-perfect technologies may coexist with human decision makers for long periods, making the design of hybrid oversight regimes a first-order managerial problem. Future work can build on these findings by experimentally varying key implementation features, such as observability of overrides, symmetry of remedies, and whether overrides affect evaluations and careers, to map when AI oversight shifts the frontier inward versus when it primarily reshapes behavior through incentives and social costs.

The main goal of this paper is to leverage a unique empirical setting to study what the adoption of AI oversight in tennis can reveal about human behavior. We do not aim to suggest an optimal practice for this application, as it is obvious that, in terms of accuracy alone, line umpires in tennis are strictly dominated by Hawk-Eye. Before the introduction of Hawk-Eye review, points were incorrectly awarded 13.9% of the time for close calls. After the introduction of Hawk-Eye review, points were incorrectly awarded 9% of the time. A system where Hawk-Eye was the decision-maker would eliminate such mistakes almost entirely. However, the ATP only replaced line umpires in the majority of top tournaments in 2025. This provides an important lesson about labor displacement: if it took 20 years and a pandemic to replace line umpires with a superior (and almost perfect) technology, it is hard to imagine AI replacing doctors or judges anytime soon. Collaboration between AI and humans will be a topic of first-order relevance for years to come, and we believe this paper provides valuable information to guide the design of better AI interventions.

We view our paper as an initial building block for understanding the implications of AI oversight on humans. Future research could be conducted in experimental settings to complement our findings, gaining more control over AI and human costs, and exploring in more detail the underlying mechanisms.

References

- Abramitzky, Ran, Liran Einav, Shimon Kolkowitz, and Roy Mill (2012). “On the Optimality of Line Call Challenges in Professional Tennis”. *International Economic Review* 53.3, pp. 939–964.
- Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz (2023). “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology”. *Working Paper 31422, National Bureau of Economic Research*.
- Albright, Alex (2024). “The Hidden Effects of Algorithmic Recommendations”. *Working paper*.
- Almog, David (2025). “Barriers to AI Adoption: Image Concerns at Work”. *arXiv: 2511.18582*.
- Almog, David, Lucas Lippman, and Daniel Martin (2026). “When an AI Judges Your Work: The Hidden Costs of Algorithmic Assessment”. *arXiv: 2603.02076*.
- Almog, David and Daniel Martin (2024). “Rational Inattention in Games: Experimental Evidence”. *Experimental Economics* 27.4, pp. 715–742.
- Athey, Susan, Kevin Bryan, and Joshua Gans (2020). “The Allocation of Decision Authority to Human and Artificial Intelligence”. *AEA Papers and Proceedings* 110, pp. 80–84.
- Avery, Mallory, Andreas Leibbrandt, and Joseph Vecci (2023). “Does Artificial Intelligence Help or Hurt Gender Diversity? Evidence from Two Field Experiments on Recruitment in Tech”. *Working paper*.
- Balakrishnan, Maya, Kris Johnson Ferreira, and Jordan Tong (2026). “Human-Algorithm Collaboration with Private Information: Naïve Advice-Weighting Behavior and Mitigation”. *Management Science* 72.1.
- Bhattacharya, Vivek and Greg Howard (2022). “Rational Inattention in the Infield”. *American Economic Journal: Microeconomics* 14.4, pp. 348–93.
- Bobonis, Gustavo J, Luis R. Cámara Fuertes, and Rainer Schwabe (2016). “Monitoring Corruptible Politicians”. *American Economic Review* 106.8, pp. 2371–2405.
- Bolte, Lukas and Collin Raymond (2023). “Emotional Inattention”. *Working paper*.
- Boyacı, Tamer, Caner Canyakmaz, and Francis de Véricourt (2023). “Human and Machine: The Impact of Machine Input on Decision Making Under Cognitive Limitations”. *Management Science* 70.2, pp. 1258–1275.
- Brown, Zach Y and Jihye Jeon (2024). “Endogenous information and simplifying insurance choice”. *Econometrica* 92.3, pp. 881–911.
- Bursztyn, Leonardo and Robert Jensen (2017). “Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure”. *Annual Review of Economics* 9.1, pp. 131–53.
- Butera, Luigi, Robert D. Metcalfe, William Morrison, and Dmitry Taubinsky (2022). “Measuring the Welfare Effects of Shame and Pride”. *American Economic Review* 112.1, pp. 122–168.

- Camerer, Colin F. (2019). “Artificial Intelligence and Behavioral Economics”. In: *The Economics of Artificial Intelligence: An Agenda*. Ed. by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press, pp. 587–608.
- Caplin, Andrew (2023). *Science Of Mistakes, The: Lecture Notes On Economic Data Engineering*. Vol. 16. World Scientific.
- Caplin, Andrew and Mark Dean (2013). *Behavioral Implications of Rational Inattention with Shannon Entropy*. Tech. rep. National Bureau of Economic Research.
- Caplin, Andrew, David Deming, Shangwen Li, Daniel Martin, Philip Marx, Ben Weidmann, and Kadachi Ye (2025). “The ABC’s of Who Benefits from Working with AI: Ability, Beliefs, and Calibration”. *Management Science*.
- Caro, Felipe, Jean-Edouard Colliard, Elena Katok, Axel Ockenfels, Nicolas Stier-Moses, Catherine Tucker, and D. J. Wu (2025). “Introduction to the Special Issue on the Human-Algorithm Connection”. *Management Science* 72.1, pp. 1–13.
- Celsi, Richard L. and Jerry C. Olson (1988). “The Role of Involvement in Attention and Comprehension Processes”. *Journal of Consumer Research* 15.2, pp. 210–224.
- Chalfin, Aaron, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, Ludwig Jens, and Sendhil Mullainathan (2016). “Productivity and Selection of Human Capital with Machine Learning”. *American Economic Review: Papers & Proceedings* 106.5, pp. 124–127.
- Chan, David C., Matthew Gentzkow, and Chuan Yu (2022). “Selection with Variation in Diagnostic Skill: Evidence from Radiologists”. *The Quarterly Journal of Economics* 137.2, pp. 729–783.
- Cho, Sungwoo (2023). “The Effect of Robot Assistance on Skills”. *Working paper*.
- Ely, Jeffrey, Romain Gauriot, and Lionel Page (2017). “Do Agents Maximise? Risk Taking on First and Second Serves in Tennis”. *Journal of Economic Psychology* 53, pp. 135–142.
- Gauriot, Romain and Lionel Page (2019). “Does Success Breed Success? a Quasi-Experiment on Strategic Momentum in Dynamic Contests”. *The Economic Journal* 129.624, pp. 3107–3136.
- Gauriot, Romain, Lionel Page, and John Wooders (2023). “Expertise, Gender, and Equilibrium Play”. *Quantitative Economics* 14.3, pp. 981–1020.
- Goldfarb, Avi and Catherine Tucker (2011). “Online Display Advertising: Targeting and Obtrusiveness”. *Marketing Science* 30.3, pp. 389–404.
- Goldsmith-Pinkham, Paul, Peter Hull, and Michal Kolesár (2024). “Contamination Bias in Linear Regressions”. *American Economic Review* 114.12, pp. 4015–51.
- Gosnell, Greer K., John A. List, and Robert D. Metcalfe (2020). “The Impact of Management Practices on Employee Productivity: A Field Experiment with Airline Captains”. *Journal of Political Economy* 128.4, pp. 1195–1626.
- Granulo, Armin, Christoph Fuchs, and Stefano Puntoni (2019). “Psychological reactions to human versus robotic job replacement”. *Nature Human Behaviour* 3, pp. 1062–1069.

- Grimon, Marie-Pascale and Christopher Mills (2025). “Better Together?: A Field Experiment on Human-Algorithm Interaction in Child Protection”. *Working paper*.
- Hoffman, Mitchell, Lisa Kahn, and Danielle Li (2018). “Discretion in Hiring”. *The Quarterly Journal of Economics* 133.2, pp. 765–800.
- Huelden, Tobias, Vitalijs Jascisens, Lars Roemheld, and Tobias Werner (2024). “Human-Machine Interactions in Pricing: Evidence from Two Large-Scale Field Experiments”. *Working paper*.
- Iakovlev, Andrei and Annie Liang (2023). “The Value of Context: Human versus Black Box Evaluators”. *Working paper*.
- Ide, Enrique and Eduard Talamas (2025). “Artificial Intelligence in the Knowledge Economy”. *Journal of Political Economy* 133.12, pp. 3713–4112.
- Jerath, Kinshuk and Qitian Ren (2021). “Consumer Rational (In)Attention to Favorable and Unfavorable Product Information, and Firm Information Design”. *Journal of Marketing Research* 58.2, pp. 343–362.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein (2011). “Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information”. *Journal of Consumer Research* 37.5, pp. 858–873.
- Kawaguchi, Kohei (2020). “When Will Workers Follow an Algorithm? A Field Experiment with a Retail Business”. *Management Science* 67.3, pp. 1670–1695.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan (2018). “Human Decisions and Machine Predictions”. *The Quarterly Journal of Economics* 133.1, pp. 237–293.
- Kovach, Matthew, Daniel Martin, and Gerelt Tserenjigmid (2026). “Learning from an Unknown DGP: Experimental Evidence on Belief Updating with AI Recommendations”. *Working paper*.
- Kreitmeir, David and Paul A Raschky (2024). “The Heterogeneous Productivity Effects of Generative AI”. *arXiv preprint arXiv:2403.01964*.
- List, John A. (2020). “Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial”. *NBER Working Paper No. 27535*.
- Margolin, Maximilian, Marko Reimer, and Daniel Schaupp (2024). “The Effects of Real-Time Feedback on Effort and Performance: Evidence from a Natural Quasi-Experiment”. *Management Science*.
- Matějka, Filip and Alisdair McKay (2015). “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model”. *American Economic Review* 105.1, pp. 272–298.
- McLaughlin, Bryce and Jann Spiess (2024). “Algorithmic Assistance with Recommendation-Dependent Preferences”. *arXiv Preprint arXiv:2208.07626*.

- Nagin, Daniel S., James B. Rebitzer, Seth Sanders, and Lowell J. Taylor (2002). “Monitoring, Motivation, and Management: The Determinants of Opportunistic Behavior in a Field Experiment”. *American Economic Review* 92.4, pp. 850–873.
- Nielsen, Kirby and John Rehbeck (2022). “When Choices Are Mistakes”. *American Economic Review* 112.7, pp. 2237–68.
- Olken, Benjamin A. (2007). “Monitoring Corruption: Evidence from a Field Experiment in Indonesia”. *Journal of Political Economy* 115.2, pp. 200–249.
- Raghu, Maithra, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan (2019). “The Algorithmic Automation Problem: Prediction, Triage, and Human Effort”. *Working paper*.
- Rajpurkar, Pranav, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Robyn L. Ball, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng (2017). “CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning”. *Working paper*.
- Rambachan, Ashesh (2024). “Identifying Prediction Mistakes in Observational Data”. *The Quarterly Journal of Economics*.
- Raymond, Lindsey (2024). “The Market Effects of Algorithms”. *Working paper*.
- Sims, Christopher A (2003). “Implications of Rational Inattention”. *Journal of Monetary Economics* 50.3, pp. 665–690.
- Topol, Eric (2019). “High-performance medicine: the convergence of human and artificial intelligence”. *Nature Medicine* 25, pp. 44–56.
- Whitney, David, Nicole Wurnitsch, Byron Hontiveros, and Elizabeth Louie (2008). “Perceptual Mislocalization of Bouncing Balls by Professional Tennis Referees”. *Current Biology* 18, R947–9.

A Data Management

A.1 Exclusion of Clay Tournaments

The officiating dynamics in clay tournaments are different. Players are allowed to request the chair umpire to review the ball bounce marks on the surface, and changing the call based on that is permitted. The clay surface presents some particular difficulties; as the surface is continuously changing during a match, the accuracy of Hawk-Eye decreases and requires constant re-calibration to operate at its best.³⁷

Tournaments on clay surfaces rejected the idea of incorporating Hawk-Eye review. They argued that they already had a mechanism in place to review incorrect calls (ball marks) and expressed reluctance, citing concerns that Hawk-Eye might not be precise enough on this surface. Another reason for rejecting the use of Hawk-Eye on clay is to avoid controversies where the ball’s mark and Hawk-Eye might disagree. As noted in The Guardian, Lars Graf, one of the ATP’s most experienced officials, explained the decision not to use Hawk-Eye on clay: “We decided not to use Hawk-Eye on clay because it might not agree with the mark the umpire is pointing at”.³⁸

At first glance, clay tournaments may appear as an ideal control group for comparing umpire reactions in tournaments with Hawk-Eye review versus those on clay. However, several complications arise, leading us to the decision to exclude clay tournaments. Firstly, before the introduction of Hawk-Eye, umpires in clay tournaments had different incentives. They could delegate some responsibility by calling more bounces as *in* and allowing the receiving player to request a mark review if there was evidence suggesting it was *out*. Since players cannot cross to the other side of the court, the receiving player has a direct view of the ball mark, making it easier for umpires to delegate in that direction. Secondly, although clay tournaments did not allow player challenges, the camera system was installed, and Hawk-Eye predictions were shown on TV broadcasts. This could have an effect on umpires, but it should differ from corrections made in the stadium. A third reason for avoiding the use of clay data is our inability to determine from our dataset which calls were corrected after examining the clay marks.

In summary, we excluded clay tournaments due to the unique incentive scheme based on ball marks, the differing treatment they received, and the complication in observability they present.

³⁷<https://www.perfect-tennis.com/why-there-is-no-hawk-eye-on-clay/>

³⁸<https://www.theguardian.com/lifeandstyle/2009/jun/27/tennis-hawk-eye>

A.2 Merging Algorithm

Our merging algorithm was designed to identify the highest number of the 144 challenges (derived from the 43 matches with video replays) within the Hawk-Eye Base data set, while keeping the number of false positives (challenges matched into the wrong point) as low as possible. We used standard variables, such as set, game, score, distance difference, player hitting the ball, and whether it was a tie-break, in order to minimize overfitting.

Our final merging algorithm comprises eight iterations, where we start using stricter criteria and gradually relax them in subsequent iterations. Some iterations, like 5 and 6, focus on very specific situations of the match (tie-breaks), and while in the testing sample do not seem useful, they might prove useful in a bigger sample.

The algorithm merged 143 out of the 144 challenges, achieving a merge rate of 99.3%. Out of these, 136 were merged to the correct point, as validated through video auditing replays, resulting in an accuracy rate of 94.4%. It would be hard to improve this efficacy as there are many challenges with missing variables like set and game. We will now elaborate on the specific criteria used in each of the eight iterations, followed by Table A.1, which summarizes the number of challenges successfully merged (and false positives) achieved in each iteration.

Merging criteria for each iteration:

Iteration 1. Same set, game, score, and player (w/ distance difference <35 mm)

Iteration 2. Same set, game, and player (w/ distance difference <15 mm)

Iteration 3. Same set, score, and player (w/ distance difference <15 mm)

Iteration 4. Same game, score, and player (w/ distance difference <10 mm)

Iteration 5. For tie-breaks: Same game and score (w/ distance difference <35 mm)

Iteration 6. For tie-breaks: Same game. If multiple, pick closest in distance (w/ distance difference <15 mm)

Iteration 7. Same set, game and player. If multiple, pick closest point in score and then in distance (w/ distance difference <35 mm)

Iteration 8. Same set and player. If multiple, pick closest point in distance (w/ distance difference <35 mm)

Iteration	Correct Merge	False Positives	Efficacy	Challenges Left
1	98	2	98%	44
2	25	2	93%	17
3	5	0	100%	12
4	1	0	100%	11
5	0	0	—	11
6	0	0	—	11
7	5	0	100%	6
8	2	3	40%	1

Table A.1: Performance for each iteration of the merging algorithm. Out of the 144 challenges audited, only one remained unmerged.

A.3 Identification of Incorrect Calls

The Hawk-Eye Base data set contains the precise location of every ball bounce and identifies the winner of each point. However, it lacks direct information on the umpires’ calls. Therefore, we are tasked with identifying the original umpire calls, which we do through the application of four criteria. The first two determine points where the umpire incorrectly called something in, and the final two identify the other type of incorrect call (umpire calling something incorrectly out). In the period involving challenges, we first restore the umpire’s original calls following a successful challenge. Then, we assess whether any of the four criteria are satisfied. We are confident that these four criteria comprehensively identify potential mistakes in our dataset.

Criterion 1. A player’s stroke is out and wins the point

We selected strokes where the first player to hit a ball out in a given point also won the point. We identified 55 points that satisfy this criterion, and 54 of them were confirmed by the video replays to be mistakes, resulting in a 98% accuracy rate.

Criterion 2. The point continued after an out

To complement Criterion 1, we must also identify points where a player hit a ball out, and despite the error, the point continued, but the same player ultimately lost the point. To achieve this, we examine instances where, after a player hits the ball out, there are at least three more strokes registered. Out of the 25 points satisfying this criterion, 24 were umpire mistakes, yielding an accuracy rate of 96%.

We also tested a more relaxed criterion that identifies points recording at least two more strokes after an out. However, this criterion proved to be too lenient, encompassing multiple cases where the point was correctly called, and players continued hitting the ball afterward.

When the criterion is relaxed to require at least two more strokes, the accuracy rate drops to 52%. This adjustment results in the identification of one new mistake at the expense of 24 incorrectly identified instances. This is why we adhere to the 3+ rule.

Criterion 3. All the strokes of a point are in, and the player hitting last loses

We identify strokes where the last player to hit the ball in did not win the point. Out of the 25 points identified in the auditing matches, 23 were confirmed to be mistakes, resulting in a 92% accuracy rate.

Criterion 4. Observing a second serve after the first serve was in

This criterion helps identify two types of instances where there is no winner in a point as a direct consequence of an umpire mistake. First, it recognizes those first serves that bounced in but were incorrectly ruled out, resulting in no winner for the point, which then moves to a second serve. Secondly, it detects points that lasted multiple strokes but had to be replayed because the umpire stopped the point by incorrectly calling a ball out. The accuracy rate for this criterion is 86%, with 36 of the 43 selected points identified as umpire mistakes. Some inaccurately selected points by this criterion include lets (serves that hit the top of the net and the point is replayed) and other unusual events (e.g., a server touching the line while serving, and the serve not counting). The occurrence and our identification of these events should not change with the introduction of Hawk-Eye.

For this criterion, we limit the analysis to strokes within 40 mm of the line, as beyond this threshold, the criterion becomes widely inaccurate. For balls bouncing between 40-100 mm away from the line, this criterion has a 30% accuracy rate, as only seven of the 23 points audited were actual umpire mistakes.

B External Validity

To structure this section, we build around the four transparency conditions given as an *onus probandi* in List (2020).

B.1 Selection

This condition concerns the representativeness of the studied group compared to the underlying population, in particular in “observables that might impact preferences, beliefs, or individual constraints.” Professional tennis umpires, especially at the top tennis tournaments we study, have long experience and are highly skilled, so our study is more likely to generalize to other specialist populations, like expert technicians. However, these tennis umpires are not required to undertake years of formal education to be qualified for their roles, unlike doctors and judges, so to the extent that this matters for AI oversight, we might be understating the effect for such populations, which could have stronger concerns about AI oversight.

B.2 Attrition

This condition asks whether there are motivational or incentive differences between subject and target groups to maintain compliance. As noted in the introduction, one of the leading tennis officials at the time of Hawk-Eye’s introduction confirmed that many important factors remained consistent over the period we study: there were no substantial changes to the training and performance assessments of umpires, the pool of umpires, or the positioning and instructions given to line judges. This is ideal for isolating the impact of psychological factors related to AI oversight, but in practice, there may be other important factors that drive the human response to AI oversight, such as career concerns.

B.3 Naturalness

This condition addresses the artificiality of a setting, such as in the choice task. We view the task in our study as having parallels to doctors making diagnosis decisions based on images that are hard to read. However, because it is more of a perceptual task than about general intelligence, this task might not capture ego-relevant concerns, perhaps understating the effect in a general population. A potentially interesting follow-up study is to investigate how humans respond to AI oversight in ego-relevant tasks.

Another aspect of our setting is that there is a live audience of spectators. We view this

aspect of our setting as having parallels to many professional decisions that occur under public scrutiny – for example, doctors make diagnostic decisions in front of patients and colleagues, judges determine bail outcomes in court with prosecutors, clerks, and officers present, and HR decisions are often made in team settings. However, having a larger audience might induce additional human responses that do not apply to more private or less intimate settings.

A final factor about our task, which we discussed previously, is that AI is nearly perfect at the task we study. Being overruled by less precise technology may influence the behavioral forces we study in different ways, as humans can shift blame or defer shame. However, this does not mean that our findings only generalize to contexts where the overruling mechanism is highly accurate, like Hawk-Eye. Rather, what matters is that the technology in use is very precise when deployed, allowing for a wider applicability in various settings. In the Zalando category managers case, technology can block a manager’s price changes only when the algorithm detects a high probability that the decision is based on non-private information. Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) show that judges release the riskiest 1% of defendants (as predicted by the algorithm) at a rate of 48.5%, suggesting a judicial protocol in which defendants identified by the algorithm as among the riskiest are denied bail, regardless of the judge’s decision.

B.4 Scaling

This condition concerns whether the effects persist if the intervention is scaled. We have already addressed issues related to horizontal scaling (moving to other populations), but here we address issues of vertical scaling (scaled up to a larger population). For cost reasons, only the top tournaments implemented AI oversight, but all umpires at this level faced this technology, so all of the umpires’ professional peers faced the same AI oversight. Thus, our study sheds light on this level of scale. One concern could be that the widespread introduction of AI oversight might dampen the psychological costs of AI oversight. However, we see that the response to AI oversight gets larger with time, so the exposure and visibility that comes with increases in scale might actually increase the effect.

C Additional Tables and Figures




Figure C.1: The location of the umpiring crew: S = the serve-line umpire, RB = the right baseline umpire, RN = the right near long-line umpire, RC = the right center-line umpire, RF = the right far-side long-line umpire. The left side of the court has corresponding line umpires, except for the serve-line umpire, who moves to the left side when the right side player has the serve.

Tournament	Start Month	Category	Court Type	Number of Matches			
				2005	2006	2007	2008
Marseille	Feb	International (250)	Hard				25
Rotterdam	Feb	International Gold (500)	Hard			15	26
Dubai	Mar	International Gold (500)	Hard				19
Las Vegas	Mar	International (250)	Hard				19
Indian Wells	Mar	Masters (1000)	Hard	17		27	
Miami	Mar	Masters (1000)	Hard	15	18	26	
Queens	Jun	International (250)	Grass	13	19	18	
UCLA	Jul	International (250)	Hard			23	
Indianapolis	Jul	International (250)	Hard		7	23	
Washington	Jul	International (250)	Hard			21	
Montreal	Aug	Masters (1000)	Hard	20		27	
Toronto	Aug	Masters (1000)	Hard		25		
Cincinnati	Aug	Masters (1000)	Hard	11	19	24	
New Haven	Aug	International (250)	Hard			18	
Beijing	Sep	International (250)	Hard			20	
Kremlin Cup	Oct	International (250)	Carpet06 Hard07		8	15	
Madrid	Oct	Masters (1000)	Hard		30	30	
Basel	Oct	International (250)	Hard			12	
Paris Masters	Oct	Masters (1000)	Carpet06 Hard07		31	33	
Shanghai	Nov	Masters (1000)	Carpet05 Hard06-07	14	15	15	

Table C.1: Information on the 35 tournaments in the consolidated dataset. The numbers in each cell indicate the matches available in our dataset. The color represents the group of tournaments, with blue indicating Pre-HK and red indicating Post-HK.





Figure C.2: An analog of Figure 2, using a one-treatment-at-a-time approach.



(a) Serves.

(b) Non-serves.

Figure C.3: An analog of Figure 4, separating serves and non-serves.




Figure C.4: The dots represent the average rate at which umpires call a ball “in” when it bounces within a 10 mm radius (e.g., at $x=0$ is the average rate for balls bouncing between -10 mm and +10 mm). The psychometric curves were approximated using locally weighted scatterplot smoothing.




Figure C.5: Heterogeneous Impact of AI Oversight on Umpires' Performance Across 20 mm Distance Bins by Tournament Category. Matches are grouped into higher-ranked tournaments (Masters 1000) and lower-ranked tournaments (International 500 and 250). Each dot represents the coefficient on the interaction between the distance bin and PostHK, the indicator variable that equals 1 if the Hawk-Eye review system is active, analyzed separately by tournament category.




Figure C.6: Speed distribution for all the balls that bounced within 100 mm of the line, separated by non-serves and serves.