

ARTIFICIAL INTELLIGENCE CLONES

ANNIE LIANG*

January 8, 2026

ABSTRACT. Large language models, trained on personal data, are increasingly able to mimic individual personalities. These “AI clones” or “AI agents” have the potential to transform how people search for matches in contexts ranging from marriage to employment. This paper presents a theoretical framework to study the trade-off between the substantially expanded search capacity of AI representations and their imperfect representation of humans. An individual’s personality is modeled as a point in k -dimensional Euclidean space, and an individual’s AI representation is modeled as a noisy approximation of that personality. I compare two search regimes: Under in person search, each person randomly meets some number of individuals and matches to the most compatible among them; under AI-mediated search, individuals match to the person with the most compatible AI representation. I show that a finite number of in-person encounters yields a better expected match than search over infinite AI representations. Moreover, when personality is sufficiently high-dimensional, simply meeting two people in person is more effective than search on an AI platform, regardless of the size of its candidate pool.

1. INTRODUCTION

Recent advances in large language models have brought us closer to a world in which individuals are represented by “AI clones” or “AI delegates” trained on their personal data (Park et al., 2024). This technology has the potential to transform how we search and match over human candidates, particularly in settings where direct

The author thanks Keaton Ellis, Drew Fudenberg, Andrei Iakovlev, David Levine, Alexandru Lupasca, Suproteem Sarkar, and Eduard Talamas for helpful comments, and National Science Foundation Grant SES-2145352 for financial support.

*Northwestern University.

engagement is costly. Labor markets are already moving in this direction: Roughly 90% of employers already use an automated system to screen candidates (Amitabh and Ansari, 2025), with four in ten employers reporting in a survey that they intend to use AI to “talk to” candidates (Demopoulos, 2024). These changes are facilitated by new AI platforms: Replicant enables job candidates to create digital representatives of themselves (see Figure 1), while Ribbon AI and Alex enable firms to automate their job recruiter (Rocha, 2025; Temkin, 2025). A similar shift is underway for dating markets. Bumble’s founder has described a future in which “your dating concierge could go and date for you with other dating concierges” (Whittaker, 2024), and platforms such as Volar Dating and Teaser AI have already adopted business models explicitly oriented around this approach (Metz, 2024; Lin, 2024).^{1,2}



FIGURE 1. Replicant helps job candidates create AI clones to represent them in conversations with recruiters and potential employees. Source: Replicant website (Replicant, 2024).

Across these domains, the appeal of AI-mediated matching is *scale*: AI representations enable search over far more candidates than a person could feasibly meet. But representations of complex human personalities based on limited training data are necessarily imperfect. Thus, there is a tradeoff between the platform’s improved scale of search and its reduced fidelity of representation.

¹This idea also appears in the Black Mirror episode “Hang the DJ,” in which the characters are revealed to be digital copies running simulations to gauge compatibility for their real-life counterparts.

²The rich medium of unstructured conversation allows AI representations to assess compatibility across many personality traits, in contrast to traditional online dating platforms which typically represent individuals by a limited set of attributes.

This paper develops a theoretical framework to study this tradeoff. Individuals are modeled as points in k -dimensional Euclidean space, and AI clones as noisy representations of those points. I ask when search conducted over AI proxies can outperform direct human evaluation. The answer turns out to crucially depend on the dimensionality of the search space. When sufficiently many attributes are relevant to match compatibility (k is large), the value of AI-mediated search is extremely low: participating on an AI platform, no matter how large its candidate pool, is no better than meeting two people in person.

The paper then turns to the broader social implications of use of AI representations for search. Specifically, I consider an extension in which the AI platform is able to develop more accurate AI representations for individuals in a “data-rich” group compared to a “data-poor” group. Despite the two groups being otherwise identical, Subject is matched to a data-rich individual with probability exceeding one-half. This probability moreover converges to 1 as the number of dimensions grows large. These results points to a potential new form of social stratification, where an individual’s outcomes depend not just on their intrinsic characteristics, but also on how effectively artificial intelligence can understand and represent those characteristics.

Section 2 describes the model. Subject’s target match is normalized to the origin, with potential matches drawn uniformly at random from the unit ball. Subject prefers to be matched to an individual closer to their target, and Subject’s payoff is decreasing in the Euclidean distance between the target and the match. I compare two search regimes: Under in-person search, Subject randomly meets m individuals and matches to the closest among them. Under AI-mediated search, all individuals are represented by noisy perturbations of their true personality vector, henceforth their *AI representations*. Subject is matched to the individual whose AI representation is closest to the representation of Subject’s target, but Subject’s payoffs are determined by the true distance between this individual and the target.

This model is inspired by AI representations but applies broadly to search problems with two defining features. First, match quality depends on many attributes that are difficult to measure. Second, human evaluators have an intuitive and holistic understanding of match quality that does not require explicit assessment of those underlying attributes. For example, a recruiter can know that a candidate interviewed well without having precisely assessed the candidate on component dimensions. In such environments, automated search is constrained by imperfect representations of the candidates and the target, while human search is constrained by the cost of direct engagement. The central tradeoff is between the errors of automated search and the capacity constraints of human evaluation.

My results focus on a quantity I call the *AI-equivalent sample size*, which measures the value of the AI platform. The AI-equivalent sample size is the smallest m such that if Subject is permitted m draws in the in-person regime, then *no* search size advantage on the AI platform can compensate for the error introduced by the AI representations. That is, Subject would rather meet m individuals in person than search over an infinite number of individuals on the AI platform.

Section 3 presents the main results. The first result, Proposition 1, says that the AI-equivalent sample size is always finite. In other words, for any level of AI approximation error, there exists a finite number of in-person encounters sufficient to outperform any search size advantage offered by the AI platform. To show this, I demonstrate an intuitive but novel monotone-likelihood ratio property: smaller Euclidean distances between AI representations are associated with smaller Euclidean distances between the actual points. Thus the best-case match on the AI platform is an individual whose AI representation perfectly matches the representation of the target. But because these representations imperfectly represent the underlying points, the expected actual distance between them is bounded away from zero. In contrast,

under in-person search, the best of m draws will eventually approach a perfect match, and thus sufficiently many draws must outperform the AI platform.

My second main result, Theorem 1, says that when the number of dimensions of personality is large, then the AI-equivalent sample size is simply two. That is, in-person search yields a higher expected payoff so long as Subject searches over at least two people.

This theorem is a consequence of counterintuitive properties of high dimensional geometry: As the number of dimensions grows large, random points in the unit ball isolate away from one another and amass near the boundary. In the context of my model, this means that each individual becomes increasingly unique in the space of personalities. While the scarcity of good matches raises the value of searching in a large pool, the best match identified on the AI platform is increasingly driven by noise in the representation rather than true compatibility. Theorem 1 shows that the breakdown of the platform’s selection mechanism outweighs the benefits of its expanded search.

The final part of the paper considers an extension in which the AI platform has access to different quantities of data for different individuals. Individuals are categorized into two groups, a “data-rich” and a “data-poor” group, which are differentiated by how well their AI representations mimic them. Despite the groups being otherwise identical in characteristics, I show that Subject is matched to a data-rich individual with probability exceeding one-half. Moreover, this probability converges to 1 as either the number of dimensions, or the disparity in estimation errors, grows large. These results suggests that use of AI representations may systematically advantage individuals that the AI platform represents more accurately.

The remainder of the paper is organized as follows. Section 2 describes the model. Section 3 presents the main results. Section 4 presents the extension to heterogeneous individuals, and Section 5 concludes.

1.1. Technological Context. This paper is motivated by recent advances in artificial intelligence that make it possible for algorithms to function as personalized representations of specific individuals. Two developments are central to this shift. First, early evidence suggests that large language models reproduce general patterns of human conversation, and can be further adapted to a particular individual’s style when fine-tuned on their personal data. Second, unlike earlier machine-learning systems that were trained for narrowly defined prediction tasks (such as predicting medical conditions from an image scan), large language models can engage in open-ended interaction both with humans and with one another. Together, these developments make large-scale automated search over human candidates technologically plausible. They are discussed in further detail below.

1.1.1. Digital Representations of Specific Individuals. Large language models are trained on large-scale corpora of text and other media drawn from the internet and related sources. While the first wave of such models adopted generic “helpful assistant” personas,³ researchers and firms are increasingly exploring the possibility that such models can be fine-tuned on personal data—such as written text, social media posts, and recorded speech and video—to produce systems that mimic stable features of a *specific* individual’s linguistic style, preferences, and judgment patterns.

Several firms already build and market such systems, often described as “AI clones” or “digital twins” of individuals. For example, Delphi AI enables public-facing figures to deploy AI agents that interact with audiences in their distinctive style; Replicant trains AI representatives for job candidates to participate in preliminary screening or interviews on their behalf; and Personal AI allows users to use their written communications to train a system that can, for example, write emails in their voice.

These generative agents are imperfect, but there are indications that they can already approximate some stable individual preferences and decision patterns with

³There is evidence that these systems already reproduce a wide range of human tendencies, preferences, and behavioral regularities (Aher et al., 2023; Mei et al., 2024; Horton, 2023).

meaningful fidelity. For example, Park et al. (2024) find that LLM-based agents trained on short personal interviews reproduce individuals’ General Social Survey survey responses and experimental choices with accuracy close to human test–retest reliability. As models improve and richer personal data become available, it is likely that such representations will continue to improve while never achieving perfect accuracy.

1.1.2. *Open-Ended Interaction and Scalable Search.* Large language models are capable of open-ended communication both with humans and with each other. For example, AI agents can interview one another, simulate collaboration, or engage in conversation on behalf of the individuals they represent. The economic significance of these developments is that interaction itself becomes scalable. Platforms such as Ribbon AI and related ‘talent twin’ systems are using conversational agents to conduct preliminary job interviews, and dating services including Volar Dating and Teaser AI are experimenting with AI-mediated exchanges in which users’ agents converse to assess compatibility. Evaluation on these platforms is no longer constrained by human time or attention, allowing them to search across vastly larger pools of potential matches.

At the same time, digital representations are inevitably imperfect, and their fidelity may vary systematically across individuals as a function of data availability, online activity, and willingness to share personal information. Section 3 establishes baseline limitations on AI-mediated search as a consequence of this imperfect representation. Section 4 subsequently studies environments in which representation quality is heterogeneous and examines how such asymmetries shape outcomes.

1.2. **Related Literature.** This paper contributes to a growing literature on the social implications of artificial intelligence (AI), in particular to research comparing human and AI evaluation. AI predictions have been shown to outperform human experts across various prediction problems (Kleinberg et al., 2017; Rajpurkar et al.,

2017; Jung et al., 2017; Angelova et al., 2022; Agarwal et al., 2023). These papers all consider settings—such as medical diagnosis—where humans possess limited intuitive knowledge about the underlying decision problem. By contrast, the present paper is motivated by settings where human actors possess a naturally rich understanding of a complex, subjective objective. This is a novel consideration relative to the literature, and the paper arrives at a very different conclusion regarding algorithmic versus human evaluation.

There is an emerging literature on AI agents. This paper complements recent empirical work about the effectiveness of AI agents (such as Jabarian and Henkel (2025) and Sarkar (2025)) with a framework for studying the welfare implications of deploying AI agents in search.

Within economic theory, this work contributes to the rich literature on search, which has explored classic questions including how long to search for (McCall, 1970; Stigler, 1961; Weitzman, 1979), what speed to search at (Urgun and Yariv, 2024), and where to search (Callander, 2011; Malladi, 2023).⁴ The present paper focuses on a new question regarding the role of the complexity of the search space, as measured by the number of attributes (Klabjan et al., 2014). I show that dimensionality fundamentally alters optimal search behavior when search is conducted with error. This result differs from, for example, Bardhi (2023) and Malladi (2023), who show that their characterizations of optimal search *extend* but are not qualitatively changed in multiple dimensions.⁵

This paper’s analysis of the impact of complexity (as measured by the number of dimensions) relates to Ely (2011) and Fudenberg and Levine (2022). Ely (2011) shows that as an optimization space grows in complexity, decision-makers resort to

⁴In related strategic experimentation models (e.g., Bolton and Harris (1999) and Keller et al. (2005)), agents face uncertain payoffs and learn by repeatedly sampling an action—a feature absent from this paper.

⁵In Bardhi (2023) and Malladi (2023), the searcher samples without error, and the focus is on inference about an unknown mapping from attributes to payoffs.

incremental adaptations (“kludges”) which generate persistent inefficiencies. Fudenberg and Levine (2022) examine how a partially-informed agent responds to systemic shocks, and show that the optimal intervention shrinks to zero as the number of dimensions grows large. My finding that higher dimensionality imposes fundamental limits on search aligns with these papers’ high-level insights.

Key to Theorem 1 is a comparison of asymptotics as the dimensionality of the search space grows large. Rate of convergence results have a long history in economic theory, but typically involve limits in the quantity of information (Vives, 1992; Moscarini and Smith, 2002; Liang and Mu, 2020; Frick et al., 2024) or the size of a population (Harel et al., 2021; Dasaratha et al., 2023). Iakovlev and Liang (2024) consider a similar asymptotic to the present paper (namely, as the number of attributes grows large) but characterize limiting beliefs rather than the limiting value of search.

Finally, while this paper adopts the narrative of a dating platform, it diverges considerably both from the classic matching frameworks (Gale and Shapley, 1962; Roth and Sotomayor, 1992), which examine how centralized mechanisms can achieve stable or efficient outcomes for populations, and from the search and matching literature (Shimer and Smith, 2000; Chade et al., 2017), which analyzes macroeconomic outcomes and equilibrium behavior in decentralized markets with many searchers. This paper instead focuses on the decision problem of a single agent navigating a high-dimensional search space.

2. MODEL

2.1. Setting. Individuals are represented by vectors in \mathbb{R}^k , where each coordinate describes a characteristic.⁶ A focal individual, hereafter *Subject*, is searching for a match, which might be a spouse or an employee. Subject’s ideal match or *target* is

⁶Example characteristics include intellectual curiosity, emotional sensitivity, relationship to time, relationship to authority, relationship to religion, energy bandwidth, style of humor, style of decision-making, and speed of decision-making.

normalized to the origin $x_0 := (0, 0, \dots, 0)$. Potential matches $i \geq 1$ have characteristic vectors $X_i \in \mathbb{R}^k$ drawn uniformly at random from the k -dimensional unit ball $B_k \equiv \{x \in \mathbb{R}^k : \|x\| \leq 1\}$, which proxies for a first-round screen where individuals who are too far from Subject’s target are not considered. Subject’s utility from matching with individual i is

$$u(X_i) = -\|X_i - x_0\| = -\|X_i\|,$$

or the negative of the Euclidean distance between individual i and the target.

There are two search regimes. Under *in-person search*, Subject randomly meets m individuals and matches to the most compatible among them.⁷ The expected distance between the target and this individual is

$$d_k^{\text{IP}}(m) \equiv \mathbb{E} \left[\min_{1 \leq i \leq m} \|X_i\| \right].$$

Under *AI-mediated search*, individuals do not interact in person but instead participate on a platform that uses “generators” to stochastically represent the candidates and the target. Specifically, in the interaction between Subject and candidate i , Subject’s target is represented by

$$Y_{0i} = x_0 + \varepsilon_{0i},$$

and individual i is represented by

$$Y_{i0} = X_i + \varepsilon_{i0},$$

where $\varepsilon_{0i}, \varepsilon_{i0} \sim \mathcal{N}(0, \sigma^2 I_k)$ are multivariate Gaussian noise terms that are independent of each other and across i . (The notation I_k denotes the identity matrix in k dimensions, and I assume throughout that $\sigma^2 > 0$.) Figure 2 depicts this model in three dimensions, i.e., $k = 3$.

⁷The main model abstracts away from whether the other individual would also agree to such a match; see Section 3.3 for discussion.

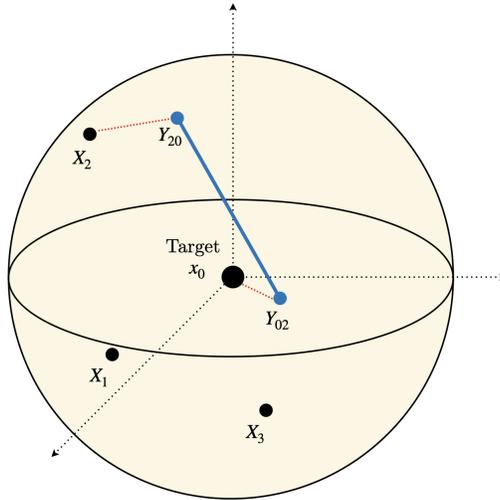


FIGURE 2. Rather than observing $\|X_2 - x_0\|$, the true distance between individual $i = 2$ and the target, the platform observes $\|Y_{20} - Y_{02}\|$, the distance between their representations.

There are n candidates on the platform. For each candidate i , the platform observes $\|Y_{0i} - Y_{i0}\|$, the distance between the representation of the target and the representation of candidate i . The platform recommends the candidate whose representation is closest to the target representation:⁸

$$i_n^* = \arg \min_{1 \leq i \leq n} \|Y_{0i} - Y_{i0}\|.$$

Crucially, although the selected candidate has the most compatible *representation*, Subject's payoffs are determined by the actual distance between Subject's target and this individual. For each sample size n , define

$$d_k^{\text{AI}}(n) \equiv \mathbb{E}[\|X_{i_n^*}\|]$$

to be the expected distance between the match i_n^* and the target, in analogy to the previous $d_k^{\text{IP}}(m)$.

⁸This is the Bayesian-optimal selection for Subject given the platform's information, i.e., $i_n^* = \arg \min_{1 \leq i \leq n} \mathbb{E}[\|X_i\| : \{\|Y_{0i} - Y_{i0}\|\}_{i=1}^n]$. See Lemma B.3 for details.

All results in this paper extend for two variations of this model. First, the representations may be fixed, so that Y_0 and Y_i are generated once and fixed in each pairwise interaction. Second, AI representations may be employed on only one side of the market, i.e., only by the searcher or by the candidates. In each of these cases, the analysis extends with minor modifications.⁹

2.2. Value of AI-Mediated Search. My main results characterize the following measure.

Definition 1. Fix any number of dimensions k . The *AI-equivalent sample size* m_k^* is the smallest $m \in \mathbb{Z}_+$ satisfying

$$d_k^{\text{IP}}(m) < d_k^{\text{AI}}(n) \quad \forall n \in \mathbb{Z}_+.$$

If the inequality above is not satisfied for any finite m , then $m_k^* = \infty$.

The AI-equivalent sample size is the smallest integer m such that searching over m individuals in person leads (in expectation) to a better match than participating on the AI platform, *no matter* the size of the platform's candidate pool. It serves as a quantitative benchmark for the value of AI-mediated search: An infinite AI-equivalent sample size means that for every m , one can choose n large enough such that search on an AI platform with n candidates outperforms in-person search over m candidates. A finite AI-equivalent sample size means that the advantage of the AI platform is fundamentally capped, and a sufficiently thorough in-person search will outperform it no matter its size. In this case, the practical value of the platform depends on the size of m_k^* .

⁹The first variation affects only the proofs of the results in Section 4, which extend when we first condition on the common confounder Y_0 . The second variation is handled by reducing the variance of the noise term when we write $Y_{i0} - Y_{0i}$ as a signal about $X_i - x_0$.

2.3. Interpretation of Model. The proposed framework applies to environments in which AI representations are already beginning to be deployed (Section 1.1), such as the following.

Example 2.1 (Hiring). A firm seeks to hire a suitable candidate. Under in-person search, a hiring manager interviews a small number of applicants. Under AI-mediated search, an AI recruiter interviews the AI representations of a substantially larger number of candidates. As discussed earlier, all results extend if an AI representation is only used on one side of this market, i.e., if an AI recruiter interviews humans or if a human recruiter interviews AI representations of candidates.

Example 2.2 (Dating). An individual seeks a partner. Under in-person search, the individual chooses the (truly) most compatible from a limited number m of candidates. Under AI-mediated search, a platform first seeks to understand the individual’s ideal match, and then evaluates a potentially much larger set of candidates $n \gg m$ relative to its understanding of the searcher’s target. In the special case where the searcher prefers someone with similar characteristics, we can interpret the target representation as a representation of the searcher themselves.

The framework also extends naturally to settings where AI representations of individuals are not yet in use but plausibly will be, such as clients seeking lawyers, founders seeking investors, or families seeking care providers.

Finally, although this paper is motivated by emerging uses of AI representations, the framework applies more broadly to search problems with two key features. First, match quality depends on many attributes that are individually difficult to measure. Second, the human searcher does not need to explicitly identify or evaluate each relevant attribute, but instead has an intuitive grasp of what constitutes a better or worse outcome.¹⁰ For example, a hiring manager can judge that an interview went poorly

¹⁰Reinforcement learning with human feedback (Christiano et al., 2017) emerged as a crucial innovation for large language models precisely because it provides a way to train systems on judgments of this kind.

without having explicitly evaluated the candidate’s communication style, problem-solving approach, or cultural fit. An AI platform, by contrast, lacks direct access to this holistic evaluation and must instead infer quality from noisy measurements of component attributes.¹¹ These features appear in other search and delegation problems as well, such as the following.

Example 2.3 (Purchasing a Home). A family seeks to purchase a home. Under in-person search, the family visits m houses directly. Alternatively, the family imperfectly communicates its preferences to a system that evaluates a larger set n of homes relative to its understanding of the target.

3. MAIN RESULTS

Section 3.1 presents two results about the limits of AI-mediated search. Section 3.2 outlines the proof of these results. Section 3.3 discusses possible extensions.

3.1. Limits of AI-Mediated Search. My first result says that the value of the AI representation regime is inherently capped: For every number of dimensions k , there is some finite number of in-person encounters that yields a better expected match than the AI platform, no matter how large its candidate pool.

Proposition 1. *For every number of dimensions $k \in \mathbb{Z}_+$, the AI-equivalent sample size m_k^* is finite.*

Intuitively, imperfect representations place a fundamental limit on the quality of search: even with an arbitrarily large platform, the expected true distance between the Subject’s target and the platform’s best-ranked match remains bounded away from zero. In contrast, under in-person search, Subject’s expected distance to the best among m individuals vanishes as m grows large. Thus for sufficiently large m , Subject’s expected match must be better in the in-person regime.

¹¹In the model, the assumption $\sigma^2 > 0$ captures this imperfect measurement and rules out settings with objectively observable attributes.

Proposition 1 holds for any number of attributes k . The next result says that when k is sufficiently large—so that many attributes are relevant to match quality—the value of AI-mediated search drops sharply.

Theorem 1. *For all k sufficiently large, the AI-equivalent sample size is $m_k^* = 2$.*

Thus, in-person search yields a better expected outcome so long as Subject searches over at least two people. The threshold dimensionality k at which this reversal occurs depends on the magnitude of AI measurement error σ^2 , but it remains finite whenever $\sigma^2 > 0$.

The implication is not that AI-mediated search is inherently ineffective, but rather that its usefulness depends critically on the effective dimensionality of match quality. When compatibility depends on a small number of characteristics, large-scale automated search can perform well. By contrast, when many distinct attributes matter, Theorem 1 implies a fundamental limitation to AI-mediated search: Even a very small amount of direct, in-person evaluation yields a better expected outcome than participation on an arbitrarily large AI platform. (This conclusion relies on the relevant attributes not being well approximated by a lower-dimensional set; otherwise, despite many nominal attributes, the effective dimensionality of evaluation would be small.)

The number of relevant attributes k likely varies across applications. In routine jobs, for instance, match quality may be predictable from a relatively small set of attributes. Other roles, such as selecting a CEO, plausibly involve a substantially larger set of traits. Long-term partner choice may involve an even richer set of relevant attributes (Joel et al., 2017, 2020)—spanning values, personality, emotional compatibility, habits, and life goals—though the effective dimensionality may also vary across individuals. Theorem 1 thus implies that AI-mediated search can be valuable when alignment on a small number of attributes suffices, but that human judgment is

indispensable—and ultimately the limiting constraint—when compatibility depends on a diverse, multidimensional set of attributes.

3.2. Proof Sketch. The proof of Theorem 1 proceeds in two main steps.

3.2.1. A uniform lower bound for $d_k^{\text{AI}}(n)$. Consider an idealized AI platform that continuously samples candidates until encountering one whose representation exactly matches the target representation, i.e., an individual i for whom $\|Y_{i0} - Y_{0i}\| = 0$. Let

$$d_k^{\text{AI}}(\infty) \equiv \mathbb{E}(\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0)$$

be the expected true distance to such a perfectly matched individual. The following lemma shows that this quantity is a uniform lower bound on $d_k^{\text{AI}}(n)$ across all n .

Lemma 1. *For all $n \in \mathbb{Z}_+$,*

$$d_k^{\text{AI}}(\infty) \leq d_k^{\text{AI}}(n).$$

To show this, I demonstrate that the pair of random variables

$$(\|X_i\|, \|Y_{i0} - Y_{0i}\|)$$

satisfies the monotone likelihood ratio property. That is, conditioning on a smaller representation distance $\|Y_{i0} - Y_{0i}\|$ shifts the conditional distribution of $\|X_i\|$ toward smaller values.¹² Consequently, the function

$$s \mapsto \mathbb{E}[\|X_i\| \mid \|Y_{i0} - Y_{0i}\| = s]$$

is increasing in s . Since the selected individual must satisfy $\|Y_{i0} - Y_{0i}\| \geq 0$, the claim follows.

¹²Given the parametric assumptions of Section 2, it is well known that the random vector $(X_i, Y_{i0} - Y_{0i})$ satisfies the monotone likelihood ratio property (MLRP), and that $(h(X_i), g(Y_{i0} - Y_{0i}))$ also satisfies MLRP for any increasing functions h and g . However, proving the specific MLRP relation for $(\|X_i\|, \|Y_{i0} - Y_{0i}\|)$ requires a novel argument because norms in \mathbb{R}^k are not monotone in each coordinate, and thus standard coordinate-wise MLRP arguments do not apply.

3.2.2. *Benchmark comparison.* Let $d_k^{\text{IP}}(1)$ denote the expected distance to a single uniformly drawn individual from the unit ball. It is well known that

$$d_k^{\text{IP}}(1) = \frac{k}{k+1}.$$

I show that both $d_k^{\text{IP}}(2)$ and $d_k^{\text{AI}}(\infty)$ converge to $d_k^{\text{IP}}(1)$ as k grows large, but crucially they do so at different rates. Specifically, the expected reduction in distance gained by picking the smaller of *two* draws rather than one is vanishing on the order of $1/k$,

$$d_k^{\text{IP}}(1) - d_k^{\text{IP}}(2) = \Theta\left(\frac{1}{k}\right)$$

while the improvement from participating on a completely saturated AI platform vanishes *strictly faster* than $1/k$,

$$d_k^{\text{IP}}(1) - d_k^{\text{AI}}(\infty) = o\left(\frac{1}{k}\right).^{13}$$

Thus $d_k^{\text{AI}}(\infty)$ exceeds $d_k^{\text{IP}}(2)$ for large k , meaning that in sufficiently high-dimensional personality spaces, Subject expects to be closer to the best of two draws in the in-person regime than to the best match among an infinite number of candidates on the AI platform. Together with Lemma 1 and an argument that $m_k^* > 1$, this yields the desired result.

3.2.3. *Intuition.* It might seem surprising that $d_k^{\text{AI}}(\infty)$ converges to $d_k^{\text{IP}}(1)$ as k grows large. The underlying force is a basic geometric feature of high-dimensional spaces. When points are drawn uniformly from the unit ball in \mathbb{R}^k , almost all probability mass concentrates in a thin shell near the boundary, and pairwise distances between points rapidly concentrate around their expectation. In the context of this model, this means that as k grows large people are increasingly unique, and compatibility with the target is increasingly hard to find. This does not mean that close points

¹³More precisely, the ratio $\frac{d_k^{\text{IP}}(1) - d_k^{\text{AI}}(\infty)}{1/k}$ vanishes to zero, while the ratio $\frac{d_k^{\text{IP}}(1) - d_k^{\text{IP}}(2)}{1/k}$ is asymptotically a positive constant.

fail to exist—on a fully saturated platform, such points appear with probability one. But the platform must identify these rare compatible candidates using noisy measurements. As a result, although $\|Y_{i0} - Y_{0i}\|$ remains informative about true compatibility in expectation (Section 3.2.1), the identity of the candidate minimizing this quantity is increasingly determined by noise rather than by genuine differences in compatibility. The platform’s ranking mechanism deteriorates, and its performance eventually converges to that of a single random draw.

3.2.4. *Simulation.* Although a bound for how large k must be for Theorem 1 to apply is beyond the scope of this paper, Figure 3 reports estimates for $d_k^{\text{IP}}(2)$ and $d_k^{\text{AI}}(N)$ with $N = 10,000$ (chosen to be some arbitrary large number) as we vary the number of dimensions k .¹⁴ The estimate of $d_k^{\text{AI}}(N)$ exceeds the estimate of $d_k^{\text{IP}}(2)$ for $k \geq 150$. Thus if match compatibility is based on 150 or more attributes, then Subject should prefer an in-person search over two individuals over participating on an AI platform with 10,000 individuals. (See Table 1 in Appendix A for the exact values in Figure 3.) Note that the small differences between $d_k^{\text{IP}}(2)$ and $d_k^{\text{AI}}(N)$ for k large do not mean that the differences between in-person search and AI-mediated search are small, since in practice we expect searchers to search over more than 2 people. (Moreover, the size of the search may be endogenously selected as a function of k .)

3.3. **Discussion.** Several aspects of the model are discussed below.

3.3.1. *The role of the parametric assumptions.* In this model, personalities are uniformly distributed within a unit ball, and AI errors follow a multivariate Gaussian distribution. The proofs of the main results do depend on these parametric assumptions—both Proposition 1 and Theorem 1 rely on the monotone-likelihood

¹⁴In more detail: in each of 1000 iterations, I draw $N = 10,000$ realizations of (X_i, Y_{i0}, Y_{0i}) in the AI representation regime and find the index i where $\|Y_{i0} - Y_{0i}\|$ is minimized. I then average over the values $\|X_{i^*}\|$ to derive an estimate of $d_k^{\text{AI}}(N)$. Likewise in the in-person regime, I draw 2 realizations of (X_1, X_2) and average over the smaller of the two norms over the 1000 trials.

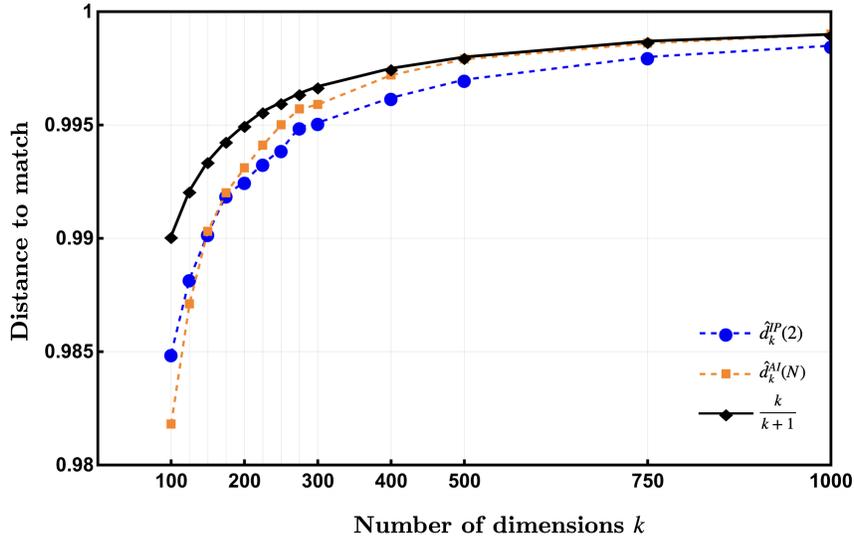


FIGURE 3. This figure reports estimates of $d_k^{IP}(2)$ and $d_k^{AI}(N)$ for $\sigma = 0.05$ and $N = 10,000$, using their average values over 1,000 draws. The quantity $d_k^{AI}(N)$ is initially below $d_k^{IP}(2)$ and eventually overtakes it. Both $d_k^{IP}(2)$ and $d_k^{AI}(N)$ converge to $k/(k+1)$ from below.

ratio property discussed in Section 3 (Lemma B.3), and Theorem 1 further uses rate-of-convergence results that rely on exact expressions. Nevertheless, the underlying mechanisms are more general. The key insight behind Proposition 1 is that the expected distance to the best match on the AI platform remains bounded away from zero (because of AI approximation errors). Meanwhile, Theorem 1 is driven by the tendency of random points to isolate from one another in high dimensions, which causes their pairwise distances to concentrate around their expected values. This phenomenon persists under a wide range of distributional assumptions.

3.3.2. *Two-sided matching.* This paper’s model of search is one-sided, and thus abstracts from the question of whether Subject’s best match (once identified) would agree to this pairing. Although I do not pursue a complete strategic model in the present paper, in many reasonable formulations of such a model, Subject’s expected payoff would decrease in the number of candidates considered by Subject’s potential

partners. This effect would further favor the in-person regime, where each individual searches over fewer candidates.

3.3.3. Increasingly accurate AI. Section 3 reveals a limitation of the AI representation regime that persists regardless of the quality of AI approximation (i.e., the size of the noise parameter σ). Nevertheless, we might expect that as the AI representations become more accurate representations of their underlying users, the AI representation regime becomes more attractive relative to the in-person regime. In the appendix, I show that for every fixed dimension k , the quantity $d_{k,\sigma}^{\text{AI}}(\infty)$ (where the noise level σ is now made explicit) converges to zero as $\sigma \rightarrow 0$. That is, Subject's expected distance to Subject's best match converges to zero as the AI's approximation errors vanish. This means that while the value to increasing the size of the AI platform is capped, the value to increasing the accuracy of the AI representations is not.

3.3.4. Sequential Search. Appendix D.3 describes a sequential search model in which the sample sizes m and n are endogenously chosen by Subject. (It is possible to mirror the initial setting by supposing that the costs of sampling in the AI regime are lower, thus micro-founding the consideration of substantially larger n on the AI platform.) I show that when there is a fixed cost to participating in the AI platform that is at least as large as the cost of searching over two people in person, then Subject prefers in-person search.

3.3.5. Human-AI Integrated Search. One might consider using the AI platform to identify m potential candidates, whom Subject then meets in person. This is unlikely to substantially improve Subject's payoffs beyond what Subject can achieve by searching over m individuals. The logic is as follows: In the proof of Theorem 1, I showed that (for large k) the best match on the AI platform is similar in distribution to a draw from a uniform distribution over the unit ball. By extension, I conjecture that (for large k) the joint distribution of the m best matches on the AI platform converges to that of m draws from a uniform distribution over the unit ball. This

would mean that for large k , Subject prefers searching over $m + 1$ people in person rather than selecting from m people sourced from the AI platform.

4. HETEROGENEOUS DATA QUANTITIES

The remainder of the paper explores heterogeneity in the quality of the AI representations. In particular, since not all individuals have the same quantity of personal data to share (e.g., because of variation in social media usage) or the same willingness to share it (e.g., because of variation in privacy preferences), the ability of the AI platform to generate representations that accurately represent each person will differ.

In this extension, individuals belong to either of two groups: a “data-rich” group for whom the approximation error is smaller, and a “data-poor” group for whom the approximation error is larger. Proposition 2 says that in this world we will see an inequality emerge between people who are better and worse understood by the AI. In particular, although the groups are otherwise identical in distribution—and thus, Subject’s actual best match is equally likely to be from either group—the probability that Subject is matched to someone in the data-rich group strictly exceeds 1/2. This probability further converges to 1 as either the number of dimensions grows large (Corollary 1) or the ratio of noise variances grows large (Corollary 2).

In more detail, the AI platform searches over n individuals in each of two groups $g \in \{R, P\}$. Individual i from group g has true personality vector X_i^g , which is independently drawn from a uniform distribution on the unit ball. When the platform simulates an interaction between Subject and this candidate, the representation of individual i is

$$Y_{i0}^g = X_i^g + \varepsilon_{i0}^g, \quad \varepsilon_{i0}^g \sim \mathcal{N}(0, \sigma_g^2 I_k)$$

and the representation of Subject’s target is

$$Y_{0i}^g = x_0 + \varepsilon_{0i}^g, \quad \varepsilon_{0i}^g \sim \mathcal{N}(0, \sigma_0^2 I_k)$$

for some constant $\sigma_0^2 > 0$.¹⁵ Crucially $\sigma_R^2 < \sigma_P^2$, so representations are more accurate for data-rich individuals. All noise terms are mutually independent.

As before, the AI platform identifies the individual whose AI representation is most compatible with Subject’s AI representation. Subject is thus matched to individual $i_{n,k}^*$ from group $g_{n,k}^*$ where

$$(i_{n,k}^*, g_{n,k}^*) = \arg \min_{g \in \{R, P\}, 1 \leq i \leq n} \|Y_{i0}^g - Y_{0i}^g\|.$$

The next result says that as the population grows large, the probability with which Subject is matched to someone from the data-rich group exceeds 1/2.

Proposition 2. *In the limit as the population grows large, the probability that Subject’s match is from the data-rich group converges to*

$$\lim_{n \rightarrow \infty} P(g_{n,k}^* = R) > \frac{1}{2}.$$

Thus, even though individuals from the data-rich group and data-poor group are identical in distribution (both are drawn uniformly at random from the unit ball), a data-rich individual is more likely to be identified as Subject’s best match.¹⁶

For a brief intuition, observe that each individual’s distance to Subject is given by $\|Y_{i0}^g - Y_{0i}^g\|$. In a large population, Subject’s best match is an individual for whom this distance is very small, i.e., an extreme draw from the underlying distribution. We thus need to analyze whether that extreme is more likely to come from the data-poor or data-rich group. I show that (a suitable transformation of) the smallest distance to any group- g individual, $M_n^g = \min_{1 \leq i \leq n} S_i^g$, is asymptotically distributed like an exponential random variable with rate proportional to the density of $Y_{i0}^g - Y_{0i}^g$ at zero—call this $f_k^g(0)$. This means that group g ’s probability of “winning,” i.e.,

¹⁵For interpretation, we might set $\sigma_0^2 \in \{\sigma_R^2, \sigma_P^2\}$, although this is not necessary for the result.

¹⁶This result is reminiscent of Cornell and Welch (1996)’s finding that more accurate signals about one group lead an employer to hire from that group with probability converging to 1 as the population grows large. But because my analysis focuses on small bounded distances rather than large unbounded qualities, the limiting probabilities here are interior.

having the smallest minimum distance to Subject, is asymptotically determined by the relative size of its exponential rate, and thus by the size of $f_k^g(0)$. Since lower noise variance leads to a more sharply peaked density at zero (i.e., $f_k^R(0) > f_k^P(0)$), the data-rich group is more likely to yield the closest representation.

The following corollaries show that this inequality is exaggerated in either of two limits: as the number of dimensions k grows large, or the disparity between noise variances σ_P^2/σ_R^2 grows large. In both cases, the probability that a data-rich individual is selected converges to 1.

Corollary 1 (Many Dimensions). *In the limit as the number of dimensions and population size both grow large, the probability that Subject’s match is from the data-rich group converges to 1, i.e.,*

$$\lim_{k \rightarrow \infty} \left(\lim_{n \rightarrow \infty} P(g_{n,k}^* = R) \right) = 1.$$

This result first takes the platform’s population size n to grow large for a fixed number of dimensions k , and then allows k to increase. This order of limits is technically convenient, since it allows us to invoke the limiting characterization in Proposition 2.¹⁷ It also has a straightforward conceptual interpretation, implying that high-dimensional geometry amplifies the large-population bias towards data-rich individuals.

The next corollary considers any sequence of noise variances $(\sigma_{P,m}^2)_{m=1}^\infty$ and $(\sigma_{R,m}^2)_{m=1}^\infty$ satisfying (1) $\lim_{m \rightarrow \infty} \sigma_{R,m}^2 = c > 0$, i.e., the smaller noise variance $\sigma_{R,m}^2$ limits to a strictly positive constant; and (2) the ratio

$$\frac{\sigma_{P,m}^2}{\sigma_{R,m}^2} \rightarrow \infty$$

¹⁷If we were to reverse these limits—taking k to infinity first—there would be no single, fixed distribution from which to draw n samples, and thus the extreme-value approach used to show Proposition 2 would not directly apply.

as $m \rightarrow \infty$, i.e., the noise variance $\sigma_{P,m}^2$ grows large relative to $\sigma_{R,m}^2$. Let $(i_{n,m}^*, g_{n,m}^*)$ denote the (random) best-match on an AI platform with n data-rich and n data-poor AI representations, whose noise variances are respectively given by $\sigma_{R,m}^2$ and $\sigma_{P,m}^2$. Corollary 2 says that as this noise ratio grows large, the large-population probability of sampling from the data-rich group converges to one.

Corollary 2 (Large Noise Disparity). *In the limit as the noise ratio and population size both grow large, the probability that Subject’s match is from the data-rich group converges to 1, i.e.,*

$$\lim_{m \rightarrow \infty} \left(\lim_{n \rightarrow \infty} P(g_{n,m}^* = R) \right) = 1.$$

In each of these results, the systematic selection of data-rich individuals benefits the Subject: the platform selects matches whose AI representations are compatible with the Subject’s target, and data-rich individuals have more accurate representations.

But this selection process has concerning distributional implications. Since the results are not specific to Subject’s characteristics, they suggest that individuals who are better understood by AI systems will be systematically more in demand. This advantage could exacerbate existing inequalities if the property of being data-rich correlates with existing group identities. They moreover point to an emerging form of social stratification: in a world of AI representations, an individual’s opportunities may depend not only on their inherent qualities but also on how effectively AI systems can capture and convey those qualities.

5. CONCLUSION

Many papers have compared human and AI evaluation on problems that are hard for both humans and machines (e.g., medical diagnosis). But there are important evaluations that are intrinsically easier for people—for example, no machine or expert knows better than ourselves whose company we enjoy. In such settings, human evaluation is more accurate than AI evaluation, but also more costly. Can automated

search over a sufficiently large number of candidates compensate for the AI's errors during evaluation? This paper suggests that scale is valuable primarily when the evaluation problem is effectively low-dimensional. If match quality can be summarized by a small number of attributes, then large-scale AI-mediated search can indeed substantially improve outcomes. In contrast, when many hard-to-measure attributes matter, expanding algorithmic scale alone cannot substitute for direct human judgment.

Several important features of AI-mediated search are abstracted away from in the present model. First, Subject's objective may be richer than Euclidean distance from a target outcome, and the platform may additionally face the problem of learning or inferring this objective from limited data. Second, both searchers and candidates may strategically manipulate or curate their AI representations. Incorporating these considerations would plausibly further limit the effectiveness of AI-mediated search, while also introducing new strategic and informational forces. They are left as interesting directions for future work.

APPENDIX A. DETAILS ON FIGURE 3

Table 1 reports the estimated values $\hat{d}_k^{\text{IP}}(2)$, $\hat{d}_k^{\text{AI}}(N)$ (for $N = 10,000$), and $\frac{k}{k+1}$ depicted in Figure 3, as well as estimates for additional values of k . For small values of k , $\hat{d}_k^{\text{AI}}(N) < \hat{d}_k^{\text{IP}}(2)$; indeed, for each of $k = 1, 5, 10$, the expected distance to the AI best match is substantially lower than the expected distance to the in-person best match. This order reverses when the number of dimensions is large (in this case, 150 or larger). Both quantities approach $k/(k+1)$ from below as k grows large.

k	$\hat{d}_k^{\text{IP}}(2)$	$\hat{d}_k^{\text{AI}}(N)$	$\frac{k}{k+1}$
1	0.3346	0.0551	0.5
5	0.7554	0.1743	0.8333
10	0.8675	0.3664	0.9091
50	0.9709	0.9126	0.9804
100	0.9849	0.9818	0.9901
125	0.9882	0.9871	0.9921
150	0.9902	0.9903	0.9934
175	0.9919	0.9920	0.9943
200	0.9925	0.9931	0.9950
225	0.9933	0.9941	0.9956
250	0.9939	0.9950	0.9960
275	0.9949	0.9957	0.9964
300	0.9951	0.9959	0.9967
400	0.9962	0.9972	0.9975
500	0.9970	0.9979	0.9980
750	0.9980	0.9986	0.9987
1000	0.9985	0.9990	0.9990

TABLE 1. This table reports estimates of $d_k^{\text{IP}}(2)$ and $d_k^{\text{AI}}(N)$ for $\sigma = 0.05$ and $N = 10,000$. The cell with the smallest value is highlighted in gray.

APPENDIX B. PROOFS OF THE MAIN RESULTS

Appendix B.1 and Appendix B.2 present key supporting results about search in the in-person and AI representation regimes. These results are used in Appendix B.3 to prove Proposition 1, and in Appendix B.4 to prove Theorem 1.

B.1. Supporting Results about In-Person Search. Lemma B.1 provides a closed-form expression for $d_k^{IP}(m)$, the expected distance between Subject's target and Subject's match under in-person search with m samples. Lemma B.2 uses this lemma to bound the distance between $d_k^{IP}(2)$ and the benchmark $d_k^{IP}(1) = k/(k+1)$.

Lemma B.1. *For every number of dimensions k and sample size m , the expected distance between Subject and Subject's match in the in-person regime is*

$$d_k^{IP}(m) = E \left[\min_{1 \leq i \leq m} \|X_i\| \right] = \frac{1}{k} B \left(\frac{1}{k}, m+1 \right)$$

where B denotes the Beta function.

Proof. First observe that

$$\begin{aligned} P \left(\min_{1 \leq i \leq m} \|X_i\| > r \right) &= P(\|X_1\| > r)^m && \text{since } X_i \text{ are i.i.d.} \\ &= \left(1 - \frac{V_{r,k}}{V_{1,k}} \right)^m = (1 - r^k)^m \end{aligned}$$

where $V_{r,k} = \frac{\pi^{k/2} r^k}{\Gamma(\frac{k}{2}+1)}$ is the volume of a ball with radius r . Thus

$$\begin{aligned} E \left[\min_{1 \leq i \leq m} \|X_i\| \right] &= \int_0^1 P \left(\min_{1 \leq i \leq m} \|X_i\| > r \right) dr \\ &= \int_0^1 (1 - r^k)^m dr \\ &= \frac{1}{k} B \left(\frac{1}{k}, m+1 \right) \end{aligned}$$

as desired. □

Lemma B.2. *For every number of dimensions k ,*

$$(B.1) \quad d_k^{IP}(2) = d_k^{IP}(1) - \frac{k}{(2k+1)(k+1)} = \frac{k}{k+1} - \frac{k}{(2k+1)(k+1)}.$$

Proof. Recalling that Lemma B.1 implies $d_k^{\text{IP}}(2) = \frac{1}{k}B\left(\frac{1}{k}, 3\right)$, let us first show

$$(B.2) \quad \frac{1}{k}B\left(\frac{1}{k}, 3\right) = \frac{k}{k+1} - \frac{k}{(2k+1)(k+1)}.$$

We can rewrite

$$\begin{aligned} \frac{1}{k}B\left(\frac{1}{k}, 3\right) &= \frac{1}{k} \frac{\Gamma\left(\frac{1}{k}\right)\Gamma(3)}{\Gamma\left(\frac{1}{k}+3\right)} && \text{by definition of the Beta function} \\ &= \frac{1}{k} \frac{2\Gamma\left(\frac{1}{k}\right)}{\Gamma\left(\frac{1}{k}+3\right)} && \text{since } \Gamma(3) = 2! \\ &= \frac{1}{k} \frac{2\Gamma\left(\frac{1}{k}\right)}{\left(\frac{1}{k}+2\right)\left(\frac{1}{k}+1\right)\frac{1}{k}\Gamma\left(\frac{1}{k}\right)} && \text{by the Gamma function recurrence relation} \\ &= \frac{2k^2}{(2k+1)(k+1)} \end{aligned}$$

Since

$$\frac{k}{k+1} - \frac{k}{(2k+1)(k+1)} = \frac{2k^2}{(2k+1)(k+1)}$$

we have the desired identity in (B.2). By similar arguments,

$$\frac{1}{k}B\left(\frac{1}{k}, 2\right) = \frac{1}{k} \frac{\Gamma\left(\frac{1}{k}\right)\Gamma(2)}{\Gamma\left(\frac{1}{k}+2\right)} = \frac{1}{k} \frac{\Gamma\left(\frac{1}{k}\right)}{\Gamma\left(\frac{1}{k}+2\right)} = \frac{1}{k} \frac{\Gamma\left(\frac{1}{k}\right)}{\left(\frac{1}{k}+1\right)\frac{1}{k}\Gamma\left(\frac{1}{k}\right)} = \frac{k}{k+1}$$

so $d_k^{\text{IP}}(2) = d_k^{\text{IP}}(1) - \frac{k}{(2k+1)(k+1)}$. That is, in expectation the second in-person draw reduces the distance between Subject and Subject's match by $\frac{k}{(2k+1)(k+1)}$. \square

B.2. Supporting Results about AI-Mediated Search. This section reports the following results: First, the expected distance between Subject's target and individual i is increasing in the distance between their AI representations (Lemma B.3).¹⁸ Thus for every finite n , Subject's expected payoff can be upper bounded by Subject's payoff in an idealized setting where some AI representation perfectly matches the representation of the target, henceforth denoted $d_k^{\text{AI}}(\infty)$ (Corollary B.1). Lemma B.4 proves that $d_k^{\text{AI}}(\infty)$ is equal to the expectation of the norm of a multivariate

¹⁸Lemma D.1 uses this to show that Subject's expected payoff given n samples is monotonically increasing in n .

Gaussian vector, conditional on that norm being less than 1. Finally, Proposition B.5 asymptotically bounds the difference between $d_k^{\text{AI}}(\infty)$ and the benchmark distance $k/(k+1)$, and Lemma B.6 shows that $d_k^{\text{AI}}(\infty) < k/(k+1)$ for every k .

Lemma B.3. $\mathbb{E}[\|X_i\| : \|Y_{i0} - Y_{0i}\| = s]$ is increasing in s .

Proof. Fix an arbitrary individual i and define $Z_i \equiv \varepsilon_{i0} - \varepsilon_{0i} \sim \mathcal{N}(0, 2\sigma^2 I_k)$. Let $R \equiv \|X\|$ and $S \equiv \|X + Z\|$, where I drop the i subscript throughout. I will show that the joint density $g_{R,S}$ satisfies the monotone likelihood ratio property, i.e.,

$$(B.3) \quad g_{R,S}(r, s) g_{R,S}(r', s') > g_{R,S}(r, s') g_{R,S}(r', s)$$

for every $1 \geq r > r' \geq 0$ and $s > s' \geq 0$. It is well known that (B.3) implies that $E(R \mid S = s)$ is increasing in s (see e.g., Milgrom (1981)).

Towards demonstrating (B.3), I will first derive a closed-form expression for each of $g_R(r)$ and $g_{S|R=r}(s)$. Since X is uniformly distributed on the unit ball, its norm R has distribution function

$$G_R(r) = \mathbb{P}(R \leq r) = \frac{V_{r,k}}{V_{1,k}} = r^k \quad \forall r \in [0, 1]$$

where $V_{r,k} = \frac{\pi^{k/2} r^k}{\Gamma(\frac{k}{2}+1)}$ is the volume of a ball with radius r . Differentiating yields

$$(B.4) \quad g_R(r) = kr^{k-1} \quad \forall r \in [0, 1].$$

Next turn to $g_{S|R=r}$. Conditional on $R = r$, we can write $X = ru$ for some unit vector u on the k -dimensional unit sphere. Hence

$$S = \|X + \varepsilon\| = \|ru + \varepsilon\|$$

where the direction u is irrelevant for the distribution of S .

Define $T = \frac{S}{\sigma}$. Then $T \mid (R = r) = \|\frac{ru}{\sigma} + \frac{\varepsilon}{\sigma}\|$ has a noncentral chi distribution with dimension k and mean vector $\mu = \frac{ru}{\sigma}$. The pdf for such a distribution is

$$(B.5) \quad h_T(t) = \frac{1}{\lambda^\nu} t^{\nu+1} \exp\left[-\frac{t^2 + \lambda^2}{2}\right] I_\nu(\lambda t)$$

where $\lambda = \|\mu\| = \frac{r}{\sigma}$, $\nu = \frac{k}{2} - 1$, and I_ν is the modified Bessel function of the first kind. Applying the change of variable $S = \sigma T$ we have

$$g_{S|R=r}(s) = \frac{1}{\sigma} h_T\left(\frac{s}{\sigma}\right)$$

for all $s \geq 0$. Substituting from (B.5) with $t = s/\sigma$, $\lambda = r/\sigma$, and $\nu = \frac{k}{2} - 1$ yields

$$(B.6) \quad g_{S|R=r}(s) = \frac{1}{\sigma^{\nu+2}} \left(\frac{s^{\nu+1}}{r^\nu}\right) \exp\left[-\frac{r^2 + s^2}{2\sigma^2}\right] I_\nu\left(\frac{rs}{\sigma^2}\right).$$

Together with (B.4), the joint density of (R, S) is therefore

$$(B.7) \quad \begin{aligned} g_{R,S}(r, s) &= g_R(r) g_{S|R=r}(s) \\ &= \frac{k}{\sigma^{k/2+1}} r^{k/2} s^{k/2-1} \exp\left[-\frac{r^2 + s^2}{2\sigma^2}\right] I_{\frac{k}{2}-1}\left(\frac{rs}{\sigma^2}\right). \end{aligned}$$

A sufficient condition for (B.3) is that $g_{R,S}$ is log-supermodular (Athey, 2002).

Define

$$(B.8) \quad \begin{aligned} \phi(r, s) &\equiv \ln g_{R,S}(r, s) \\ &= \ln \frac{k}{\sigma^{k/2+1}} + \frac{k}{2} \ln r + \left(\frac{k}{2} - 1\right) \ln s - \frac{r^2 + s^2}{2\sigma^2} + \ln I_{\frac{k}{2}-1}\left(\frac{rs}{\sigma^2}\right). \end{aligned}$$

Since only the final term of (B.8) has a nonzero cross-partial,

$$\frac{\partial^2}{\partial r \partial s} \phi(r, s) = \frac{\partial^2}{\partial r \partial s} \ln \left(I_{\frac{k}{2}}\left(\frac{rs}{\sigma^2}\right) \right)$$

Let $z = rs/\sigma^2$ and again denote $\nu = \frac{k}{2} - 1$. Then

$$\frac{\partial}{\partial r} \ln I_\nu(z) = \frac{I'_\nu(z)}{I_\nu(z)} \cdot \frac{s}{\sigma^2}$$

and so

$$\begin{aligned} \frac{\partial^2}{\partial r \partial s} \ln \left(I_\nu \left(\frac{rs}{\sigma^2} \right) \right) &= \frac{\partial}{\partial s} \left[\frac{I'_\nu(z)}{I_\nu(z)} \cdot \frac{s}{\sigma^2} \right] = \frac{\partial}{\partial s} \left[\frac{I'_\nu(z)}{I_\nu(z)} \right] \cdot \frac{s}{\sigma^2} + \frac{I'_\nu(z)}{I_\nu(z)} \cdot \frac{1}{\sigma^2} \\ &= \frac{\partial}{\partial z} \left[\frac{I'_\nu(z)}{I_\nu(z)} \right] \cdot \frac{rs}{\sigma^4} + \frac{I'_\nu(z)}{I_\nu(z)} \cdot \frac{1}{\sigma^2} \end{aligned}$$

We can now sign the RHS using properties of I_ν , the modified Bessel function of the first kind (see e.g., Árpád Baricz (2010)). First, $I'_\nu(z) \geq 0$ and $I_\nu(z) \geq 0$ for all $z \geq 0$ (which is implied by $r, s \geq 0$); thus, $\frac{I'_\nu(z)}{I_\nu(z)} \geq 0$. Moreover, for $\nu \geq -1/2$ (which is satisfied in our setting since $k \geq 1$), the function I_ν is log-convex on $(0, \infty)$, i.e., $\frac{d}{dz} \left[\frac{I'_\nu(z)}{I_\nu(z)} \right] \geq 0$. Thus the entire RHS is positive, so $g(r, s) = \exp(\phi(r, s))$ is log-supermodular as desired. \square

Corollary B.1. *Define $d_k^{AI}(\infty) = E[\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0]$. Then $d_k^{AI}(\infty) \leq d_k^{AI}(n)$ for all $n \in \mathbb{Z}_+$.*

Proof. Define $W = \min_{1 \leq j \leq n} \|Y_{j0} - Y_{0j}\|$. Then

$$d_k^{AI}(n) = E[\|X_i\| : \|Y_{i0} - Y_{0i}\| = W]$$

where clearly $P(W \geq 0) = 1$. So

$$d_k^{AI}(n) \geq E[\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0]$$

follows by Lemma B.3. \square

Lemma B.4. *For each i , $E[\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0] = E[\|Z\| : \|Z\| \leq 1]$ where $Z \sim \mathcal{N}(0, 2\sigma^2 I_k)$ and I_k is the identity matrix in k dimensions.*

Proof. Recall that $Y_{i0} = Y_{0i}$ is equivalent to $X_i + \varepsilon_{i0} = x_0 + \varepsilon_{0i}$, or more simply

$$X_i = Z_i$$

where $Z_i = \varepsilon_{0i} - \varepsilon_{i0} \sim \mathcal{N}(0, 2\sigma^2 I_k)$. By Bayes' rule, the posterior density of X_i conditional on the event $\{X_i = Z_i\}$ satisfies

$$f_{X_i|Z_i=X_i}(x) = \frac{f_{X_i}(x)f_{Z_i}(x)}{\int_{x \in B_k} f_{X_i}(x)f_{Z_i}(x)dx}$$

Since $f_{X_i}(x)$ is constant on B_k and zero elsewhere, this further simplifies to

$$f_{X_i|Z_i=X_i}(x) = \frac{f_{Z_i}(x)}{\int_{x \in B_k} f_{Z_i}(x)dx}$$

That is, $X_i | (Y_{i0} = Y_{0i}) \stackrel{d}{=} Z_i | \{\|Z_i\| \leq 1\}$ which further implies

$$E[\|X_i\| : Y_{i0} = Y_{0i}] = E[\|Z_i\| : \|Z_i\| \leq 1]$$

as desired. □

Lemma B.5. $d_k^{AI}(\infty) = \frac{k}{k+1} + o\left(\frac{1}{k}\right)$

Proof. Let $Z \sim N(0, 2\sigma^2 I_k)$ and $R = \|Z\|$. Since $W = \frac{R^2}{2\sigma^2} \sim \chi_k^2$ has density

$$f_W(w) = \frac{1}{2^{k/2}\Gamma(\frac{k}{2})} w^{\frac{k}{2}-1} e^{-w}$$

the change of variables $r = \sqrt{2\sigma^2 w}$ yields

$$(B.9) \quad f_R(r) = \frac{2^{1-k}}{\Gamma(\frac{k}{2})\sigma^k} r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) \quad \forall r \geq 0.$$

Therefore

$$\mathbb{E}[R | R \leq 1] = \frac{\int_0^1 r f_R(r) dr}{\int_0^1 f_R(r) dr} = \frac{\int_0^1 r^k \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}{\int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}.$$

Define

$$I_1(k) := \int_0^1 r^k \exp\left(-\frac{r^2}{4\sigma^2}\right) dr, \quad I_2(k) := \int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr.$$

We will show

$$(B.10) \quad \frac{I_1(k)}{I_2(k)} = \frac{k}{k+1} + O\left(\frac{1}{k^2}\right)$$

Define $g(r) := \exp\left(-\frac{r^2}{4\sigma^2}\right)$ on the domain $[0, 1]$. Note that $g \in C^2([0, 1])$ and

$$g(1) = e^{-\frac{1}{4\sigma^2}}, \quad g'(r) = -\frac{r}{2\sigma^2} e^{-\frac{r^2}{4\sigma^2}}, \quad g'(1) = -\frac{1}{2\sigma^2} e^{-\frac{1}{4\sigma^2}}.$$

Moreover, g'' is continuous on $[0, 1]$, so $M := \sup_{r \in [0, 1]} |g''(r)| < \infty$. For each integer $k \geq 1$, integrate by parts with $u = g(r)$ and $dv = r^k dr$:

$$\begin{aligned} \int_0^1 r^k g(r) dr &= \left[\frac{r^{k+1}}{k+1} g(r) \right]_0^1 - \frac{1}{k+1} \int_0^1 r^{k+1} g'(r) dr \\ &= \frac{g(1)}{k+1} - \frac{1}{k+1} \int_0^1 r^{k+1} g'(r) dr. \end{aligned}$$

Apply integration by parts again to the remaining integral, with $u = g'(r)$ and $dv = r^{k+1} dr$:

$$\begin{aligned} \int_0^1 r^{k+1} g'(r) dr &= \left[\frac{r^{k+2}}{k+2} g'(r) \right]_0^1 - \frac{1}{k+2} \int_0^1 r^{k+2} g''(r) dr \\ &= \frac{g'(1)}{k+2} - \frac{1}{k+2} \int_0^1 r^{k+2} g''(r) dr. \end{aligned}$$

Combining these expressions yields

$$\int_0^1 r^k g(r) dr = \frac{g(1)}{k+1} - \frac{g'(1)}{(k+1)(k+2)} + \frac{1}{(k+1)(k+2)} \int_0^1 r^{k+2} g''(r) dr.$$

The remainder is bounded as

$$\left| \int_0^1 r^{k+2} g''(r) dr \right| \leq M \int_0^1 r^{k+2} dr = \frac{M}{k+3},$$

so

$$(B.11) \quad \int_0^1 r^k g(r) dr = \frac{g(1)}{k+1} - \frac{g'(1)}{(k+1)(k+2)} + O\left(\frac{1}{k^3}\right).$$

Thus

$$I_1(k) = \frac{g(1)}{k+1} - \frac{g'(1)}{(k+1)(k+2)} + O\left(\frac{1}{k^3}\right)$$

$$I_2(k) = \frac{g(1)}{k} - \frac{g'(1)}{k(k+1)} + O\left(\frac{1}{k^3}\right)$$

Factor out $g(1) = e^{-1/(4\sigma^2)}$ and define $A := -\frac{g'(1)}{g(1)} = \frac{1}{2\sigma^2}$. Then

$$I_1(k) = g(1) \left(\frac{1}{k+1} + \frac{A}{(k+1)(k+2)} + O\left(\frac{1}{k^3}\right) \right)$$

$$I_2(k) = g(1) \left(\frac{1}{k} + \frac{A}{k(k+1)} + O\left(\frac{1}{k^3}\right) \right)$$

Their ratio simplifies to

$$\frac{I_1(k)}{I_2(k)} = \frac{k + \frac{Ak}{k+2} + O\left(\frac{1}{k}\right)}{(k+1) + A + O\left(\frac{1}{k}\right)}.$$

Since $\frac{Ak}{k+2} = A + O(1/k)$, the numerator is $k + A + O(1/k)$ and the denominator is $k + 1 + A + O(1/k)$, hence

$$\frac{I_1(k)}{I_2(k)} = \frac{k + A + O\left(\frac{1}{k}\right)}{k + 1 + A + O\left(\frac{1}{k}\right)} = \frac{k}{k+1} + O\left(\frac{1}{k^2}\right).$$

This establishes (B.10), and therefore

$$\mathbb{E}[R \mid R \leq 1] = \frac{I_1(k)}{I_2(k)} = \frac{k}{k+1} + o\left(\frac{1}{k}\right),$$

as desired. □

Lemma B.6. *For every positive integer k , $d_k^{AI}(\infty) < \frac{k}{k+1}$.*

Proof. Let T be the random variable with density function

$$g_T(r) = \frac{r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right)}{\int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}$$

on $[0, 1]$, observing that

$$E[T] = d_k^{\text{AI}}(\infty) = \frac{\int_0^1 r^k \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}{\int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}.$$

Consider also a Beta($k, 1$) random variable B , whose density function on $[0, 1]$ is $g_B(r) = kr^{k-1}$ and expectation is $E[B] = \frac{k}{k+1}$. The likelihood ratio

$$\frac{g_T(r)}{g_B(r)} = \frac{1}{kr^{k-1}} \frac{r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right)}{\int_0^1 r^{k-1} \exp(-r^2/4\sigma^2) dr} = \frac{1}{k} \frac{\exp(-r^2/4\sigma^2)}{\int_0^1 r^{k-1} \exp(-r^2/4\sigma^2) dr}.$$

is strictly decreasing in r , implying that $\left\{\frac{g_T(r)}{g_B(r)}\right\}_{r \geq 0}$ has the monotone likelihood ratio property. Thus the distribution of B first-order stochastically dominates the distribution of T , implying in particular that $E[T] < E[B]$, or equivalently,

$$\frac{\int_0^1 r^k \exp\left(-\frac{r^2}{4\sigma^2}\right) dr}{\int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr} < \frac{k}{k+1}$$

as desired. □

B.3. Proof of Proposition 1. By Lemma B.1,

$$d_k^{\text{IP}}(m) = \frac{1}{k} B\left(\frac{1}{k}, m+1\right)$$

For every fixed k , this expression converges to zero as m grows large. Since the quantity $\mathbb{E}[\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0]$ is strictly positive, we can identify a finite m such that

$$d_k^{\text{IP}}(m) < \mathbb{E}[\|X_i\| : \|Y_{i0} - Y_{0i}\| = 0].$$

Then by Corollary B.1,

$$d_k^{\text{IP}}(m) < d_k^{\text{AI}}(n) \quad \forall n \in \mathbb{Z}_+$$

implying that the AI equivalent sample size exists and is finite.

B.4. Proof of Theorem 1. By Lemma B.2,

$$d_k^{\text{IP}}(2) = \frac{k}{k+1} - \frac{k}{(2k+1)(k+1)}$$

where $\frac{k}{(2k+1)(k+1)} = \Theta\left(\frac{1}{k}\right)$. That is, $d_k^{\text{IP}}(2)$ is smaller than $\frac{k}{k+1}$ by a quantity that is on the order of $1/k$. By Lemma B.5,

$$d_k^{\text{AI}}(\infty) = \frac{k}{k+1} + o\left(\frac{1}{k}\right)$$

i.e., the difference between $d_k^{\text{AI}}(\infty)$ and $k/(k+1)$ vanishes faster than $1/k$ as k grows large. Moreover by Corollary B.1, $d_k^{\text{AI}}(\infty)$ is a uniform lower bound on $d_k^{\text{AI}}(n)$ for all n . Thus when k is sufficiently large, we have

$$d_k^{\text{IP}}(2) < d_k^{\text{AI}}(n)$$

for all $n \in \mathbb{Z}_+$, so $m_k^* \leq 2$. Since also

$$d_k^{\text{IP}}(1) = \mathbb{E}[\|X_1\|] = d_k^{\text{AI}}(1)$$

we have $m_k^* > 1$. Thus $m_k^* = 2$ is the AI equivalent sample size for k sufficiently large.

APPENDIX C. PROOFS OF RESULTS IN SECTION 4

C.1. Proof of Proposition 2. For each individual i in group g , define

$$S_i^g = \|Y_{i0}^g - Y_{0i}^g\| = \|X_i^g + Z_i^g\| \quad \text{where} \quad Z_i^g = \varepsilon_{i0}^g - \varepsilon_{0i}^g$$

to be the distance between this individual's AI clone Y_{i0}^g and Subject's AI clone Y_{0i}^g in their interaction. Recall that X_i^g is drawn from a uniform distribution on the unit ball while $Z_i^g \sim \mathcal{N}(0, \nu_g I_k)$, where $\nu_R \equiv \sigma_R^2 + \sigma_0^2$ and $\nu_P \equiv \sigma_P^2 + \sigma_0^2$. So the random variable

$$Y_{i0}^g - Y_{0i}^g = X_i^g + Z_i^g$$

admits a density f_k^g that is locally Lipschitz around the origin. That is, there is an $L > 0$ such that

$$|f_k^g(x) - f_k^g(0)| \leq L\|x\|$$

for all x sufficiently close to the origin. Integrating this Lipschitz bound over the ball of small radius s , we obtain

$$\begin{aligned} \left| \int_{\|x\| \leq s} (f_k^g(x) - f_k^g(0)) dx \right| &\leq \int_{\|x\| \leq s} |f_k^g(x) - f_k^g(0)| dx \\ &\leq \int_{\|x\| \leq s} L\|x\| dx = O(s^{k+1}) \end{aligned}$$

Thus

$$\int_{\|x\| \leq s} f_k^g(x) dx = \int_{\|x\| \leq s} f_k^g(0) dx + O(s^{k+1})$$

or equivalently,

$$(C.1) \quad H_g(s) = f_k^g(0)V_{1,k}s^k + O(s^{k+1})$$

where H_g denotes the CDF of S_i^g , and $V_{1,k} = \frac{\pi^{k/2}}{\Gamma(\frac{k}{2}+1)}$ is the volume of the unit ball.

Now let $M_n^g = \min_{1 \leq i \leq n} S_i^g$ be the smallest distance to Subject among individuals in group g . Since S_g^i are iid,

$$P(M_n^g > s) = (1 - H_g(s))^n,$$

and plugging in (C.1) yields

$$P(M_n^g > s) = \left(1 - f_k^g(0)V_{1,k}s^k + O(s^{k+1})\right)^n.$$

Consider an arbitrary $r > 0$ and set $s_n = \left(\frac{r}{n}\right)^{1/k}$. Then as n grows large,

$$P(M_n^g > s_n) = \left(1 - f_k^g(0)V_{1,k}\left(\frac{r}{n}\right) + O\left(\left(\frac{1}{n}\right)^{(k+1)/k}\right)\right)^n \rightarrow e^{-\lambda_g r}$$

where $\lambda_g \equiv f_k^g(0)V_{1,k}$. If we define the scaled random variable $Y_n^g = n(M_n^g)^k$, then equivalently

$$P(Y_n^g > r) \longrightarrow e^{-\lambda_g r} \quad \text{as } n \text{ grows large}$$

So $Y_n^g \rightarrow Y^g$ in distribution, where the limiting random variable Y^g follows an exponential distribution with rate parameter λ_g . Since moreover (Y_n^R, Y_n^P) are independent for each n , and the limiting random variables (Y^R, Y^P) are independent,

$$P(Y_n^R < Y_n^P) \longrightarrow P(Y^R < Y^P) = \frac{\lambda_R}{\lambda_R + \lambda_P}$$

where the final equality is a well-known comparison for the hit rate of two independent exponential random variables. Since the latter expression reduces to $\frac{f_k^R(0)}{f_k^R(0)+f_k^P(0)}$, and the event $\{Y_n^R < Y_n^P\}$ is identical to the event $\{M_n^R < M_n^P\}$, we have the first part of the desired result.

It remains to show that $f_k^R(0) > f_k^P(0)$. Using the parametric distributions of X_i^g and Z_g^i , we can analytically derive the density at zero to be

$$(C.2) \quad f_k^g(0) = \int_{x \in B_k} \frac{1}{V_{1,k}} \frac{1}{(2\pi\nu_g)^{k/2}} \exp\left(-\frac{\|x\|^2}{2\nu_g}\right) dx > 0$$

Moving to spherical coordinates,

$$\int_{x \in B_k} \exp\left(-\frac{\|x\|^2}{2\nu_g}\right) dx = |S^{k-1}| \int_0^1 r^{k-1} \exp\left(-\frac{r^2}{2\nu_g}\right) dr$$

where $|S^{k-1}| = \frac{2\pi^{k/2}}{\Gamma(\frac{k}{2})}$ is the surface area of the unit sphere. Further performing the change of variable $t = \frac{r^2}{2\nu_g}$ yields

$$\int_{x \in B_k} \exp\left(-\frac{\|x\|^2}{2\nu_g}\right) dx = |S^{k-1}| \frac{(2\nu_g)^{k/2}}{2} \int_0^{\frac{1}{2\nu_g}} t^{\frac{k}{2}-1} e^{-t} dt$$

Plugging this into (C.2) and noting that $\Gamma(\frac{k}{2} + 1) = \frac{k}{2}\Gamma(\frac{k}{2})$, we obtain

$$(C.3) \quad f_k^g(0) = \frac{k}{2\pi^{k/2}} \int_0^{\frac{1}{2\nu_g}} t^{\frac{k}{2}-1} e^{-t} dt$$

which is clearly monotonically decreasing in ν_g . Thus $\nu_R < \nu_P$ implies $f_k^R(0) > f_k^P(0)$, as desired.

C.2. Proof of Corollary 1. From Proposition 2, we know that the limiting probability as the population grows large is

$$\lim_{n \rightarrow \infty} P(g_{n,k}^* = R) = \frac{f_k^R(0)}{f_k^R(0) + f_k^P(0)}.$$

Moreover by (C.3),

$$\frac{f_k^R(0)}{f_k^P(0)} = \frac{\int_0^{\frac{1}{2\nu_R}} t^{\frac{k}{2}-1} e^{-t} dt}{\int_0^{\frac{1}{2\nu_P}} t^{\frac{k}{2}-1} e^{-t} dt}$$

Observe that

$$e^{-\frac{1}{2\nu_g}} \frac{2}{k} \left(\frac{1}{2\nu_g} \right)^{\frac{k}{2}} \leq \int_0^{\frac{1}{2\nu_g}} t^{\frac{k}{2}-1} e^{-t} dt \leq \frac{2}{k} \left(\frac{1}{2\nu_g} \right)^{\frac{k}{2}}$$

since $e^{-\frac{1}{2\nu_g}} \leq e^{-t} \leq 1$ on the domain $t \in [0, 1/(2\nu_g)]$. Thus

$$\frac{\int_0^{\frac{1}{2\nu_R}} t^{\frac{k}{2}-1} e^{-t} dt}{\int_0^{\frac{1}{2\nu_P}} t^{\frac{k}{2}-1} e^{-t} dt} \geq \frac{e^{-\frac{1}{2\nu_R}} \frac{2}{k} \left(\frac{1}{2\nu_R} \right)^{\frac{k}{2}}}{\frac{2}{k} \left(\frac{1}{2\nu_P} \right)^{\frac{k}{2}}} = e^{-\frac{1}{2\nu_R}} \left(\frac{\nu_P}{\nu_R} \right)^{\frac{k}{2}}$$

which converges to ∞ as k grows large, since $\nu_P/\nu_R = (\sigma_0^2 + \sigma_P^2)/(\sigma_0^2 + \sigma_R^2) > 1$ by assumption that $\sigma_R^2 < \sigma_P^2$.

C.3. Proof of Corollary 2. By (C.2),

$$\frac{f_k^R(0)}{f_k^P(0)} = \left(\frac{\sigma_{0,m}^2 + \sigma_{P,m}^2}{\sigma_{0,m}^2 + \sigma_{R,m}^2} \right)^{k/2} \cdot \frac{\int_{B_k} \exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{R,m}^2)}\right) dx}{\int_{B_k} \exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{P,m}^2)}\right) dx}.$$

By assumption that $\sigma_{P,m}^2/\sigma_{R,m}^2 \rightarrow \infty$ as $m \rightarrow \infty$ while $\sigma_{0,m}^2$ is fixed, we have

$$\left(\frac{\sigma_{0,m}^2 + \sigma_{P,m}^2}{\sigma_{0,m}^2 + \sigma_{R,m}^2} \right)^{k/2} \rightarrow \infty$$

as $m \rightarrow \infty$. Next let

$$D_m \equiv \int_{B_k} \exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{P,m}^2)}\right) dx$$

be the integral in the denominator. Since $\sigma_{P,m}^2 \rightarrow \infty$ and $\sigma_{0,m}^2 \geq 0$, we have $\sigma_{0,m}^2 + \sigma_{P,m}^2 \rightarrow \infty$. Hence, for each fixed $x \in B_k$,

$$\exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{P,m}^2)}\right) \rightarrow 1$$

as $m \rightarrow \infty$. Moreover, for all m and all $x \in B_k$,

$$0 \leq \exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{P,m}^2)}\right) \leq 1,$$

and the constant function 1 is integrable over the bounded set B_k . Thus by the Dominated Convergence Theorem,

$$D_m \rightarrow \int_{B_k} 1 dx = V_{1,k},$$

where $V_{1,k}$ denotes the volume of the unit ball in \mathbb{R}^k . In particular, D_m is bounded above and bounded away from zero for all sufficiently large m .

Now define

$$N_m \equiv \int_{B_k} \exp\left(-\frac{\|x\|^2}{2(\sigma_{0,m}^2 + \sigma_{R,m}^2)}\right) dx$$

to be the integral in the numerator. By assumption,

$$\lim_{m \rightarrow \infty} N_m = c > 0,$$

and hence $\liminf_{m \rightarrow \infty} N_m \geq c > 0$.

Combining these observations,

$$\liminf_{m \rightarrow \infty} \frac{N_m}{D_m} = \frac{\liminf_{m \rightarrow \infty} N_m}{\lim_{m \rightarrow \infty} D_m} \geq \frac{c}{V_{1,k}} > 0.$$

Therefore,

$$\frac{f_k^R(0)}{f_k^P(0)} = \left(\frac{\sigma_{0,m}^2 + \sigma_{P,m}^2}{\sigma_{0,m}^2 + \sigma_{R,m}^2} \right)^{k/2} \cdot \frac{N_m}{D_m} \rightarrow \infty$$

as $m \rightarrow \infty$. The desired result then follows from Proposition 2.

APPENDIX D. ADDITIONAL RESULTS

Section D.1 proves a result that says that Subject's expected payoff in the AI representation regime is monotonically increasing in n , the number of candidates. Section D.2 proves a result that says that the expected distance between Subject and Subject's AI best match (in the infinite sample benchmark) converges to zero as $\sigma \rightarrow 0$.

D.1. Monotonicity in n .

Lemma D.1. *Subject's expected payoff in the AI representation regime is monotonically increasing in n ; that is, $d_k^{AI}(n+1) < d_k^{AI}(n)$ for every $n \in \mathbb{Z}_+$.*

Proof. For each i define

$$S_i \equiv \|Y_{i0} - Y_{0i}\| = \|X_i + \varepsilon_{i0} - \varepsilon_{0i}\| = \|X_i + Z_i\|$$

where $Z_i \sim \mathcal{N}(0, 2\sigma^2 I_k)$. Consider a single probability space on which the infinite sequence $\{(X_i, S_i)\}_{i=1}^\infty$ is defined, noting that this tuple is independent across i . On this space, the random variable

$$i_n^* = \arg \min_{1 \leq i \leq n} S_i$$

is well-defined for every n . Moreover we can write

$$d_k^{AI}(n+1) = \mathbb{E} [\|X_{n+1}\| \cdot \mathbb{1}(S_{n+1} < S_{i_n^*}) + \|X_{i_n^*}\| \cdot \mathbb{1}(S_{n+1} \geq S_{i_n^*})]$$

$$d_k^{AI}(n) = \mathbb{E} [\|X_{i_n^*}\| \cdot \mathbb{1}(S_{n+1} < S_{i_n^*}) + \|X_{i_n^*}\| \cdot \mathbb{1}(S_{n+1} \geq S_{i_n^*})]$$

so (by linearity of expectation) it is sufficient to show

$$\mathbb{E} [\|X_{n+1}\| \cdot \mathbb{1}(S_{n+1} < S_{i_n^*})] \leq \mathbb{E} [\|X_{i_n^*}\| \cdot \mathbb{1}(S_{n+1} < S_{i_n^*})]$$

or equivalently

$$\mathbb{E} [\|X_{n+1}\| : S_{n+1} < S_{i_n^*}] \leq \mathbb{E} [\|X_{i_n^*}\| : S_{n+1} < S_{i_n^*}]$$

Conditional on S_{n+1} , the variable X_{n+1} is independent of the sequence $\{(X_i, S_i)\}_{i=1}^n$. Since $S_{i_n^*}$ is a measurable function of $\{(X_i, S_i)\}_{i=1}^n$, also $X_{n+1} \perp\!\!\!\perp S_{i_n^*} \mid S_{n+1}$. Thus

$$(D.1) \quad E [\|X_{n+1}\| : S_{n+1}, S_{i_n^*}] = \mathbb{E} [\|X_{n+1}\| : S_{n+1}] =: m(S_{n+1})$$

So also

$$\begin{aligned} E [\|X_{n+1}\| : S_{n+1} < S_{i_n^*}] &= E [E [\|X_{n+1}\| : S_{n+1}, S_{i_n^*}] : S_{n+1} < S_{i_n^*}] && \text{by L.I.E.} \\ &= E [m(S_{n+1}) \mid S_{n+1} < S_{i_n^*}] && \text{by (D.1)} \end{aligned}$$

Similarly,

$$(D.2) \quad E [\|X_{i_n^*}\| : S_{n+1}, S_{i_n^*}] = \mathbb{E} [\|X_{i_n^*}\| : S_{i_n^*}] =: m(S_{i_n^*})$$

and

$$\begin{aligned} E [\|X_{i_n^*}\| : S_{n+1} < S_{i_n^*}] &= E [E [X_{i_n^*} \mid S_{n+1}, S_{i_n^*}] \mid S_{n+1} < S_{i_n^*}] \\ &= E [m(S_{i_n^*}) \mid S_{n+1} < S_{i_n^*}] \end{aligned}$$

Finally, observe that on the event $\{S_{n+1} < S_{i_n^*}\}$ we have $S_{n+1}(\omega) < S_{i_n^*}(\omega)$ pointwise for every $\omega \in \{S_{n+1} < S_{i_n^*}\}$. Since by Lemma B.3 the function $m(\cdot)$ is monotonically increasing,

$$m(S_{n+1}(\omega)) \leq m(S_{i_n^*}(\omega))$$

also holds pointwise on $\{S_{n+1} < S_{i_n^*}\}$. This inequality is thus preserved by taking an expectation on $\{S_{n+1} < S_{i_n^*}\}$, i.e.,

$$E [m(S_{n+1}) \mid S_{n+1} < S_{i_n^*}] \leq E [m(S_{i_n^*}(\omega) \mid S_{n+1} < S_{i_n^*})].$$

Finally, by (D.1) and (D.2), this is equivalent to the desired statement

$$E [\|X_{n+1}\| : S_{n+1} < S_{i_n^*}] \leq E [\|X_{i_n^*}\| : S_{n+1} < S_{i_n^*}].$$

□

D.2. Limit as $\sigma \rightarrow 0$. Define

$$I_1(\sigma) := \int_0^1 r^k \exp\left(-\frac{r^2}{4\sigma^2}\right) dr \quad \text{and} \quad I_2(\sigma) := \int_0^1 r^{k-1} \exp\left(-\frac{r^2}{4\sigma^2}\right) dr.$$

Apply the change of variable $t = \frac{r^2}{4\sigma^2}$ to obtain

$$I_1(\sigma) = 2^k \sigma^{k+1} \int_0^{1/(4\sigma^2)} t^{\frac{k-1}{2}} e^{-t} dt = 2^k \sigma^{k+1} \gamma\left(\frac{k+1}{2}, \frac{1}{4\sigma^2}\right),$$

where $\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt$ is the lower incomplete gamma function. The same substitution for $I_2(\sigma)$ yields

$$I_2(\sigma) = 2^{k-1} \sigma^k \int_0^{1/(4\sigma^2)} t^{\frac{k}{2}-1} e^{-t} dt = 2^{k-1} \sigma^k \gamma\left(\frac{k}{2}, \frac{1}{4\sigma^2}\right).$$

From the above identities, we have

$$\frac{I_1(\sigma)}{I_2(\sigma)} = 2\sigma \frac{\gamma\left(\frac{k+1}{2}, \frac{1}{4\sigma^2}\right)}{\gamma\left(\frac{k}{2}, \frac{1}{4\sigma^2}\right)}.$$

Since for each fixed $s > 0$, $\lim_{x \rightarrow \infty} \gamma(s, x) = \Gamma(s)$ (where Γ represents the Gamma function),

$$\lim_{\sigma \rightarrow 0} \gamma\left(\frac{k+1}{2}, \frac{1}{4\sigma^2}\right) = \Gamma\left(\frac{k+1}{2}\right) \quad \text{and} \quad \lim_{\sigma \rightarrow 0} \gamma\left(\frac{k}{2}, \frac{1}{4\sigma^2}\right) = \Gamma\left(\frac{k}{2}\right)$$

Thus $\lim_{\sigma \rightarrow 0} \frac{I_1(\sigma)}{I_2(\sigma)} = 0 \cdot \frac{\Gamma\left(\frac{k+1}{2}\right)}{\Gamma\left(\frac{k}{2}\right)} = 0$, as desired.

D.3. Sequential Search. Time $t = 0, 1, 2, \dots$ is discrete. In each period $t \geq 1$, an individual X_t is drawn from a uniform distribution on the unit ball B_k , and error terms ε_{t0} and ε_{0t} are drawn from $\mathcal{N}(0, \sigma^2 I_k)$. All random variables are mutually independent. The state space is $\Omega = (B_k)^\infty \times \mathbb{R}^\infty \times \mathbb{R}^\infty$ with typical element $\omega = ((X_t)_{t=1}^\infty, (\varepsilon_{t0})_{t=1}^\infty, (\varepsilon_{0t})_{t=1}^\infty)$.

At time $t = 0$, Subject chooses between search in the in-person regime and the AI representation regime. If Subject chooses the in-person regime, then in each subsequent period t , Subject observes X_t (corresponding to meeting individual t) and decides whether to stop and match with the closest individual so far or to continue searching. At the time of stopping τ , Subject receives the payoff

$$- \min\{\|X_1\|, \|X_2\|, \dots, \|X_\tau\|\} - c_{\text{IP}}(\tau),$$

where the in-person search cost $c_{\text{IP}}(\cdot)$ is an increasing function of τ .

If instead Subject chooses the AI representation regime, then Subject first pays a fixed cost κ to use the AI-based platform. In each period t , Subject observes the representation match quality

$$S_t = \|X_t + \varepsilon_{t0} - \varepsilon_{0t}\|$$

and decides whether to stop or continue. If Subject stops at time τ , the realized payoff is

$$-\|X_{i^*(S_{\leq \tau})}\| - c_{\text{AI}}(\tau) - \kappa,$$

where $S_{\leq t} = (S_1, \dots, S_t)$ is Subject's history at time t , and $i^*(S_{\leq \tau}) = \arg \min_{1 \leq i \leq \tau} S_i$ denotes the individual prior to time τ whose AI representation is closest to Subject's AI representation. The AI platform search cost $c_{\text{AI}}(\cdot)$ is an increasing function of τ . (It is reasonable to assume that $c_{\text{AI}}(\tau) \leq c_{\text{IP}}(\tau)$ for all τ , but this is not necessary for the results.)

Subject's strategy thus consists of an action $a \in \{\text{IP}, \text{AI}\}$ and a stopping rule τ_a , which is a map $\Omega \rightarrow [0, \infty)$ that is measurable with respect to the observed histories in the chosen regime. Formally, let $\mathcal{F}_t^{\text{IP}}$ denote the filtration generated by the in-person history $X_{\leq t} = (X_1, \dots, X_t)$, and let $\mathcal{F}_t^{\text{AI}}$ denote the filtration generated by the AI history $S_{\leq t}$. Then if $a = \text{IP}$, the stopping rule must satisfy $\{\tau_a \leq t\} \in \mathcal{F}_t^{\text{IP}}$ for every t , and if $a = \text{AI}$, the stopping rule must satisfy $\{\tau_a \leq t\} \in \mathcal{F}_t^{\text{AI}}$ for every t .

The following result says that when the number of dimensions is sufficiently large, and if the fixed cost of joining the AI platform exceeds the cost of searching over two individuals in person, then Subject optimally chooses the in-person regime.

Proposition D.1. *Suppose $\kappa > c_{\text{IP}}(2)$. Then for all k sufficiently large, Subject optimally chooses $a = \text{IP}$ at period $t = 0$.*

Proof. Since (X_t, S_t) are i.i.d. across t ,

$$E [\|X_{i^*(S_{\leq t})}\| : S_{i^*(S_{\leq t})}] = E [\|X_{i^*(S_{\leq t})}\| : S_{\leq t}].$$

We have already shown in Lemma B.3 that $E [\|X_{i^*(S_{\leq t})}\| | S_{i^*(S_{\leq t})} = s]$ is increasing in s , and thus

$$E [\|X_{i^*(S_{\leq t})}\| : S_{i^*(S_{\leq t})} = s] \geq d_k^{\text{AI}}(\infty) = E [\|X_{i^*(S_{\leq t})}\| : S_{i^*(S_{\leq t})} = 0]$$

So, pointwise for each realization of ω , we have

$$-E [\|X_{i^*(\omega)}(\omega)\| : S_{i^*(\omega)}(\omega)] \leq -d_k^{\text{AI}}(\infty)$$

where $i^*(\omega) \equiv i^*(S_{\leq \tau}(\omega))$. This inequality is preserved by integrating over ω ; thus

$$-E [\|X_{i^*(\omega)}(\omega)\|] = -E [E [\|X_{i^*(\omega)}(\omega)\| : S_{i^*(\omega)}(\omega)]] \leq -d_k^{\text{AI}}(\infty)$$

So for every stopping rule τ_{AI} in the AI representation regime, we have

$$-\mathbb{E}_{\tau_{\text{AI}}} [\|X_{i^*(S_{\leq \tau_{\text{AI}}})}\| + c_{\text{AI}}(\tau_{\text{AI}}) + \kappa] \leq -d_k^{\text{AI}}(\infty) - \kappa$$

The proof of Theorem 1 established $d_k^{\text{AI}}(\infty) \geq d_k^{\text{IP}}(2)$ for k sufficiently large. Moreover, $\kappa \geq c_{\text{IP}}(2)$ by assumption. Thus

$$-d_k^{\text{AI}}(\infty) - \kappa \leq -d_k^{\text{IP}}(2) - c_{\text{IP}}(2) \quad \text{for } k \text{ sufficiently large}$$

Finally observe that

$$-d_k^{\text{IP}}(2) - c_{\text{IP}}(2) \leq \sup_{\tau_{\text{IP}}} \mathbb{E} [-\min\{\|X_1\|, \|X_2\|, \dots, \|X_{\tau_{\text{IP}}}\|\} - c_{\text{IP}}(\tau_{\text{IP}})]$$

since one feasible stopping rule τ_{IP} is to sample twice and stop regardless of the realized history. Thus we have the desired result. \square

REFERENCES

- AGARWAL, N., A. MOEHRING, P. RAJPURKAR, AND T. SALZ (2023): “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology,” Working Paper 31422, National Bureau of Economic Research.
- AHER, G. V., R. I. ARRIAGA, AND A. T. KALAI (2023): “Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies,” in *Proceedings of the 40th International Conference on Machine Learning*, ed. by A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, PMLR, vol. 202 of *Proceedings of Machine Learning Research*, 337–371.
- AMITABH, U. AND A. ANSARI (2025): “Hiring with AI doesn’t have to be so inhumane,” *World Economic Forum*, accessed 7 Jan 2026.
- ANGELOVA, V., W. DOBBIE, , AND C. S. YANG (2022): “Algorithmic Recommendations and Human Discretion,” Working Paper.
- ATHEY, S. (2002): “Monotone Comparative Statics under Uncertainty,” *The Quarterly Journal of Economics*, 117, 187–223.
- BARDHI, A. (2023): “Attributes: Selective Learning and Influence,” Working Paper.

- BOLTON, P. AND C. HARRIS (1999): “Strategic Experimentation,” *Econometrica*, 67, 349–374.
- CALLANDER, S. (2011): “Searching for Good Policies,” *American Political Science Review*, 105, 643–662.
- CHADE, H., J. EECKHOUT, AND L. SMITH (2017): “Sorting through Search and Matching Models in Economics,” *Journal of Economic Literature*, 55, 493–544.
- CHRISTIANO, P. F., J. LEIKE, T. B. BROWN, M. MARTIC, S. LEGG, AND D. AMODEI (2017): “Deep Reinforcement Learning from Human Preferences,” in *Advances in Neural Information Processing Systems (NIPS)*, 4302–4310.
- CORNELL, B. AND I. WELCH (1996): “Culture, Information, and Screening Discrimination,” *Journal of Political Economy*, 104, 542–571.
- DASARATHA, K., B. GOLUB, AND N. HAK (2023): “Learning from Neighbours about a Changing State,” *Review of Economic Studies*, 90, 2326–2369.
- DEMOPOULOS, A. (2024): “The job applicants shut out by AI: ‘The interviewer sounded like Siri’,” *The Guardian*, accessed 7 Jan 2026.
- ELY, J. C. (2011): “Kludged,” *American Economic Journal: Microeconomics*, 3, 210–231.
- FRICK, M., R. IJIMA, AND Y. ISHII (2024): “Monitoring with Rich Data,” Working paper.
- FUDENBERG, D. AND D. K. LEVINE (2022): “Intervention with limited information,” *International Journal of Game Theory*, 51, 379–389, published online 19 October 2021.
- GALE, D. AND L. S. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *The American Mathematical Monthly*, 69, 9–15.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): “Rational Groupthink,” *The Quarterly Journal of Economics*, 136, 621–668, published: 08 July 2020.

- HORTON, J. J. (2023): “Large Language Models as Simulated Economic Agents: What Can We Learn from *Homo Silicus*?” NBER Working Paper 31122, National Bureau of Economic Research.
- IAKOVLEV, A. AND A. LIANG (2024): “The Value of Context: Human versus Black Box Evaluators,” Working Paper.
- JABARIAN, B. AND L. HENKEL (2025): “Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews,” Working paper.
- JOEL, S., P. W. EASTWICK, C. J. ALLISON, X. B. ARRIAGA, Z. G. BAKER, E. BARR-KALIF, AND ET AL. (2020): “Machine learning uncovers the most robust self-report predictors of relationship quality across 43 longitudinal couples studies,” *Proceedings of the National Academy of Sciences*, 117, 19061–19071.
- JOEL, S., P. W. EASTWICK, AND E. J. FINKEL (2017): “Is Romantic Desire Predictable? Machine Learning Applied to Initial Romantic Attraction,” *Psychological Science*, 28, 1478–1489.
- JUNG, J., C. CONCANNON, R. SHROFF, S. GOEL, AND D. G. GOLDSTEIN (2017): “Simple rules for complex decisions,” Working Paper.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73, 39–68.
- KLABJAN, D., W. OLSZEWSKI, AND A. WOLINSKY (2014): “Attributes,” *Games and Economic Behavior*, 88, 190–206.
- KLEINBERG, J., H. LAKKARAJU, J. LESKOVEC, J. LUDWIG, AND S. MULLAINATHAN (2017): “Human Decisions and Machine Predictions,” *The Quarterly Journal of Economics*, 133, 237–293.
- LIANG, A. AND X. MU (2020): “Complementary Information and Learning Traps*,” *The Quarterly Journal of Economics*, 135, 389–448.
- LIN, B. (2024): “Grindr Aims to Build the Dating World’s First AI ‘Wingman’,” *The Wall Street Journal*, accessed: 2025-02-15.

- MALLADI, S. (2023): “Searching in the Dark and Learning Where to Look,” Working paper, draft date: April 19, 2023.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 84, 113–126.
- MEI, Q., Y. XIE, W. YUAN, AND M. O. JACKSON (2024): “A Turing test of whether AI chatbots are behaviorally similar to humans,” *Proceedings of the National Academy of Sciences*, 121, e2313925121.
- METZ, C. (2024): “AI Is Changing Dating Apps, But Not in the Way You Think,” *The New York Times*.
- MILGROM, P. R. (1981): “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, 12, 380–391.
- MOSCARINI, G. AND L. SMITH (2002): “The Law of Large Demand for Information,” *Econometrica*, 70, 2351–2366.
- PARK, J. S., C. Q. ZOU, A. SHAW, B. M. HILL, C. CAI, M. R. MORRIS, R. WILLER, P. LIANG, AND M. S. BERNSTEIN (2024): “Generative Agent Simulations of 1,000 People,” *arXiv*, accessed 2025-01-01.
- RAJPURKAR, P., J. IRVIN, K. ZHU, B. YANG, H. MEHTA, T. DUAN, D. DING, A. BAGUL, C. LANGLOTZ, K. SHPANSKAYA, M. P. LUNGREN, AND A. Y. NG (2017): “CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning,” Working Paper.
- REPLICANT (2024): “About Replicant: Revolutionizing recruitment through AI-powered candidate replication,” <https://replicant.im/about>, accessed 2025-01-01.
- ROCHA, N. (2025): “Welcome to Your Job Interview. Your Interviewer Is A.I.” *The New York Times*, accessed 2025-01-06.
- ROTH, A. E. AND M. SOTOMAYOR (1992): “Two-Sided Matching,” in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 1, 485–541.

- SARKAR, S. K. (2025): “AI Agents, Productivity, and Higher-Order Thinking: Early Evidence from Software Development,” Preprint.
- SHIMER, R. AND L. SMITH (2000): “Assortative Matching and Search,” *Econometrica*, 68, 343–369.
- STIGLER, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69, 213–225.
- TEMKIN, M. (2025): “AI recruiter Alex raises \$17M to automate initial job interviews,” Accessed January 7, 2026.
- URGUN, C. AND L. YARIV (2024): “Contiguous Search: Exploration and Ambition on Uncharted Terrain,” *Journal of Political Economy*, 132, published online December 18, 2024.
- VIVES, X. (1992): “How Fast do Rational Agents Learn?” *Review of Economic Studies*, 60, 329–347.
- WEITZMAN, M. L. (1979): “Optimal Search for the Best Alternative,” *Econometrica*, 47, 641–654.
- WHITTAKER, R. (2024): “Would you let AI date for you? Bumble’s founder thinks that could be the future,” <https://www.forbes.com.au/news/innovation/bumble-ai-dating-concierge/>, forbes Australia, accessed January 2026.
- ÁRPÁD BARICZ (2010): “Bounds for Modified Bessel Functions of the First and Second Kinds,” *Proceedings of the Edinburgh Mathematical Society*, 53, 575–599.