



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Learning to Successfully Hire in Online Labor Markets

Marios Kokkodis, Sam Ransbotham

To cite this article:

Marios Kokkodis, Sam Ransbotham (2023) Learning to Successfully Hire in Online Labor Markets. Management Science 69(3):1597-1614. <https://doi.org/10.1287/mnsc.2022.4426>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Learning to Successfully Hire in Online Labor Markets

Marios Kokkodis,^{a,*} Sam Ransbotham^a

^aCarroll School of Management, Boston College, Chestnut Hill, Massachusetts 02467

*Corresponding author

Contact: kokkodis@bc.edu,  <https://orcid.org/0000-0002-5037-6060> (MK); sam.ransbotham@bc.edu,

 <https://orcid.org/0000-0001-5305-035X> (SR)

Received: March 22, 2019

Revised: November 29, 2020; February 17, 2022

Accepted: March 22, 2022

Published Online in Articles in Advance:
August 26, 2022

<https://doi.org/10.1287/mnsc.2022.4426>

Copyright: © 2022 INFORMS

Abstract. Hiring in online labor markets involves considerable uncertainty: which hiring choices are more likely to yield successful outcomes and how do employers adjust their hiring behaviors to make such choices? We argue that employers will initially explore the value of available information. When employers observe successful outcomes, they will keep reinforcing their hiring strategies; but when the outcomes are unsuccessful, employers will adjust their hiring behaviors. To investigate these dynamics, we propose a two-component framework that links hiring choices with task outcomes. The framework's first component, a Hidden Markov Model, captures how employers transition from unsuccessful to successful hiring decisions. The framework's second component, a conditional logit model, estimates employer hiring choices. Analysis of 238,364 hiring decisions from a large online labor market shows that, often, employers initially explore cheaper contractors with a lower reputation. When these options result in unsuccessful outcomes, employers learn and adjust their hiring behaviors to rely more on reputable contractors who are not as cheap. Such hiring tends to be successful, guiding employers to reinforce their hiring processes. As a result, the market observes employers transition from cheaper, lower-reputation options with poorer performance to more expensive reputable options with better performance. We attribute part of this behavior to employer confidence and risk attitude, which can change over time. This work is the first to investigate how employers learn to make successful hiring choices in online labor markets. As a result, it provides platform managers with new knowledge and analytics tools to target employer interventions.

History: Accepted by Anandhi Bharadwaj, information systems.

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/mnsc.2022.4426>.

Keywords: employer evolution • successful hiring choices • online labor markets • empirical analysis

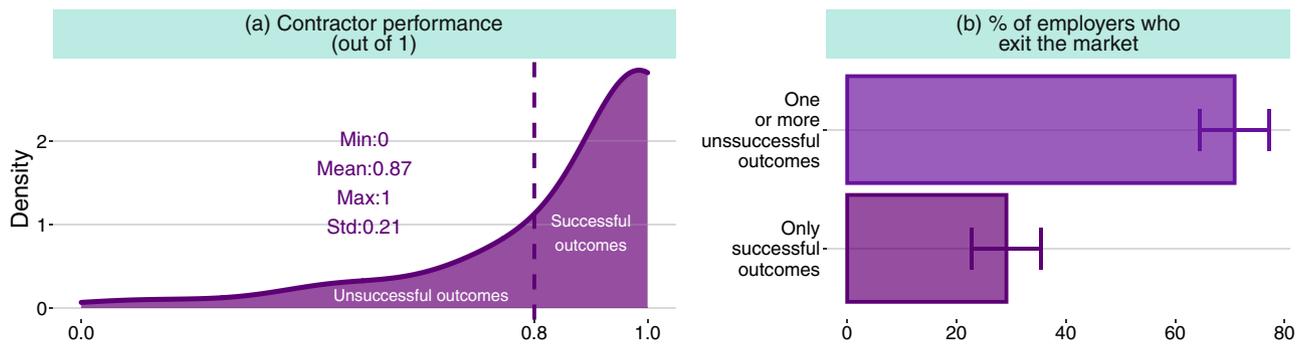
1. Introduction

Online platforms overcome many market inefficiencies by facilitating transactions between consumers and producers throughout the world. Online labor markets, such as Upwork and Freelancer, are prime examples; they connect independent employers with contractors who accomplish diverse tasks. These platforms boomed over the past decade: Upwork, for example, now has more than 14 million contractors and more than five million employers; these contractors complete three million tasks annually with a total transaction volume of over one billion U.S. dollars (Lauren 2017, Brier and Pearson 2018). This growth will likely continue (if not accelerate) in the future as the sharing economy continues to structure the future of work (Sundararajan 2016, Jain et al. 2018).

Uncertainty in these markets makes hiring difficult. Although online markets are similar to their offline counterparts, experience using offline markets does not entirely translate to these growing platforms. Available information, for example, differs between

offline and online markets. To identify the best candidate, employers in both markets must assess both observed and latent contractor characteristics. Observed characteristics include a contractor's skillset, work history, and certifications. Latent characteristics include a contractor's true knowledge and abilities (Geva and Saar-Tsechansky 2016). But employers have additional information available in online markets, particularly for signals of latent characteristics (e.g., worker reputation; Kokkodis and Ipeirotis 2016). The existence of latent characteristics and heterogeneity in the observed ones (Kokkodis and Ipeirotis 2014) creates an uncertain environment where hiring decisions rely on idiosyncratic assessments of fit between task and contractor characteristics.

This uncertainty in hiring decisions often results in unsuccessful hires that are costly for the employer—and the platform. Figure 1(a) illustrates the distribution of task outcomes (i.e., the normalized contractor performance) in an online labor market. On average, contractors perform well (mean performance is 0.87 out

Figure 1. (Color online) Uncertainty in Hiring Decisions Yields Unsuccessful and Costly Outcomes

Notes. Employers often make hiring choices that result in unsuccessful outcomes (a). Such outcomes are costly, as they discourage employers from participating in the market (b). Std, standard deviation.

of 1). However, a significant portion (27%) of hiring choices results in outcomes where contractor performance is well below 0.8. These unsuccessful outcomes are costly for the marketplace, as they discourage employers from continuing to participate on the platform (Tripp and Grégoire 2011). In fact, 71% of employers who exited the market in our data (i.e., did not post a new task for more than three consecutive months) had one or more unsuccessful outcomes; only 29% of those employers had only successful hiring choices (Figure 1(b)).

Given how costly unsuccessful hiring outcomes are, it is crucial for employers and platform managers to understand what leads to successful decisions and whether (and how) employers learn and improve their hiring behaviors. Despite extensive research on the drivers of hiring decisions in online labor markets (Yoganarasimhan 2013, Moreno and Terwiesch 2014, Pallais 2014, Gefen et al. 2016, Kokkodis and Ipeirotis 2016, Lin et al. 2016), we know very little about how these hiring decisions connect to task outcomes. For example, employers tend to hire reputable contractors (Yoganarasimhan 2013, Lin et al. 2016, Hong and Pavlou 2017) and are willing to trade off price and reputation (Moreno and Terwiesch 2014). Yet, what is not known is how these hiring preferences affect hiring outcomes and, more importantly, how employers learn to adjust their behaviors to achieve better outcomes. Hence, we ask the following:

Which hiring choices are more likely to yield successful outcomes, and how do employers adjust their hiring behaviors to make such choices?

To investigate, we focus on how employers adjust their assessment of characteristics that affect hiring outcomes in online labor markets. We argue that market participants will initially explore the value of platform information (Peterson and Pitz 1988, Busenitz and Barney 1997, Dequech 1999). While exploring, some employers might be prone to try lower-cost options that, on average, have a lower reputation (Moreno and

Terwiesch 2014). Sometimes such hires will be successful; but, on average, they will result in poorer outcomes (Kokkodis and Ipeirotis 2016). Research on individual exploitation and exploration suggests that when employers observe successful outcomes, they will keep exploiting knowledge they have developed and reinforce their hiring strategies; but when employers observe unsuccessful outcomes, they will adjust their hiring behaviors and keep exploring market information (Mom et al. 2007, Lee and Meyer-Doyle 2017, Lee 2019). Combined, these dynamics suggest that the market will observe employers moving from lower-reputation and cheaper options to higher-reputation and more expensive ones that will result in more successful outcomes.

To empirically test these mechanisms, we propose a new two-component framework that facilitates the investigation of both hiring behaviors and outcomes. In particular, the first component of the framework, a Hidden Markov Model (HMM), captures how employers transition from *unsuccessful* to *successful hiring decisions*. The second component, a conditional logit model, estimates employer hiring choices. This structure allows for a comparison of hiring behaviors as employers move between successful and unsuccessful decisions; hence, it links *hiring choices with task outcomes* for the first time.

Analysis of 238,364 hiring decisions from a large online labor market provides empirical evidence supporting the relationships we theorize. Specifically, we show that (1) employers indeed initially explore available signals and they often select lower cost and lower reputation alternatives that are more likely to yield unsuccessful outcomes and (2) employers then learn from exploration and start to exploit by reinforcing successful prior hiring behaviors. Finally, the two-component framework shows that employers who learn to rely more on contractor reputation and to not chase cheap contractors become more successful.

Why do some employers choose initially to hire lower-cost contractors with a relatively worse reputation? We

argue that confidence, defined as “belief in success,” might be a plausible mechanism. Employers with higher confidence will likely have higher risk attitudes and hence try to beat the market by hiring cheaper contractors that they consider as good deals. Indeed, some of these hires will end up being successful; but most will likely underperform (Liu et al. 2008, Danescu-Niculescu-Mizil et al. 2009, Lu et al. 2010, Kokkodis and Ipeirotis 2016, Kokkodis 2021). Our learning framework suggests that employers whose choices underperform will keep exploring and adjust their hiring behaviors and, as a result, reevaluate their belief in success and potentially adjust their confidence and risk attitudes. On the other hand, employers with lower confidence will have lower risk attitudes and will hire conservatively by selecting contractors who are not as cheap but have a good reputation. These choices will tend to be more successful and will guide many employers to reinforce (exploit) their hiring behaviors. As a result of these dynamics, some higher-confidence employers who experience unsuccessful outcomes will adjust their hiring preferences and transition to lower-risk choices that are, on average, more successful.

This research extends our understanding of how users learn in online markets, particularly in online labor markets. It is the first study to explain how hiring choices affect outcomes and how employers learn to adjust their hiring behaviors to achieve better outcomes. Through a new theoretical framework, our work argues that employers prone to making decisions that are more likely to fail will keep exploring and adjust their hiring behaviors. On the other hand, employers who hire successfully will keep reinforcing their hiring behaviors that exploit already acquired knowledge. Empirical evidence confirms this framework, showing that as employers transition from unsuccessful to successful hiring choices, they rely more on contractors’ feedback scores and become less price sensitive. Additionally, we provide some evidence that struggles on the platform could stem from employers’ beliefs in their ability to take advantage of perceived pricing anomalies but subsequently learn that the market prices and contractor reputation tend to correlate with probability of successful outcomes.

Methodologically, our work provides a new two-component framework that links hiring choices with task outcomes. This framework provides future research a way to study transactions in these markets while incorporating both employer-evolving preferences and task outcomes. Even further, our analysis helps platform managers understand how employers make successful hiring decisions. Such understanding is fundamentally important as both (1) employers who hire successfully and (2) contractors who receive positive feedback are more likely to keep participating (Jerath et al. 2011, Tripp and Grégoire 2011). In addition, our

framework allows managers to dynamically measure and track the likelihood of employers making successful hiring choices. Knowledge of such employer-evolving trends gives platform managers opportunities to target interventions and incentivize employers to adjust their hiring behaviors.

2. Theory

Online platforms in general—and online labor markets in particular (Agrawal et al. 2015, Lauren 2017, Brier and Pearson 2018)—have grown substantially over the past decades supporting the transition of many business functions from offline to online. In adopting online platforms, individuals use cognitive structures from their prior experience that allow them to form perceptions about uncertainty (Barr et al. 1992) to combine multiple sources of information and assess incomplete signal in new contexts (Peterson and Pitz 1988, Busenitz and Barney 1997, Dequech 1999). New online platform users encounter a mix of more familiar and less familiar information that they must assess to use online platforms successfully. For more familiar information, cognitive structures they formed outside the market can help their decisions in the market but may be less relevant in the new context than they realize. For less familiar information, new users may be unsure how to adapt prior cognitive structures that lack platform-specific context. As a result, new users must learn to assess market information in the new context to reduce the likelihood of unsuccessful outcomes.

2.1. Learning in Online Labor Markets

We use online labor markets as a context to study learning in such new platforms for three reasons. First, online labor markets are growing in importance (Agrawal et al. 2015, Lauren 2017, Brier and Pearson 2018). Second, users must combine observable and latent information to make a hiring decision. Third, these markets have offline parallels from which make some, but not all, platform information familiar from prior experience. Indeed, the theoretical importance of online labor markets motivated multidisciplinary research in areas such as algorithmic contractor management (Sheng et al. 2008, Horton and Chilton 2010, Ipeirotis et al. 2010, Mason and Watts 2010); design choices that increase market efficacy (Gopal et al. 2003, Arora and Forman 2007, Dey et al. 2010, Horton 2010, Allon et al. 2012, Chen and Horton 2016, Kokkodis and Ipeirotis 2016, Liang et al. 2017); and, particularly relevant for our context, employer hiring choices (Yoganarasimhan 2013, Moreno and Terwiesch 2014, Pallais 2014, Hong et al. 2015, Lin et al. 2016).

In online labor markets, potential employers post the specifications for work they need, mostly contract

work of fixed time or scope. Once posted, jobs typically attract numerous applications. Employers then base hiring decisions on expectations of contractor performance. Once hired, the contractors complete the work and the employer observes actual performance and can assess the success of the hire.

As such, labor is an “experience” good (Nelson 1970): employers do not know about a contractor’s performance on a task before the contractor completes the task. When hiring contractors, employers must consider observable (e.g., listed skills, certifications, price) and latent (actual abilities) characteristics to estimate contractor performance. The existence of latent characteristics and the heterogeneity in the observable ones (Kokkodis and Ipeirotis 2014) increases the uncertainty in hiring choices: employers must use imperfect information to make noisy estimates about contractors’ abilities.

But processing signals from such imperfect sources is subjective and correlates with an individual’s ability to acquire and process information (Grinold and Kahn 2000). Individuals can potentially learn from observing the outcomes of their decisions and adjust how they process market signals accordingly. And although prior research explains how employers choose contractors in online labor markets (Moreno and Terwiesch 2014, Chan and Wang 2017, Kokkodis 2018), it does not yet explain how (and what) employers learn to make successful hiring choices.

Specifically, prior relevant works are limited in two important ways. First, studies of hiring choices alone do not provide information about the *result* of these choices—we know a considerable amount about what contractors employers choose, but we know less about how well these choices work out. Second, static analysis of hiring choices does not reflect how (or whether)

employers adjust their hiring behaviors over time. Although it is helpful to understand employer hiring characteristics at a given time, task outcomes (successful or unsuccessful completion of a task) provide considerable valuable information to employers about the relevance of various market signals, hence helping them to reassess and learn from their prior behaviors. As a result, a study that investigates how employer learning occurs—and if it occurs—from prior outcomes helps us better understand how online labor markets work.

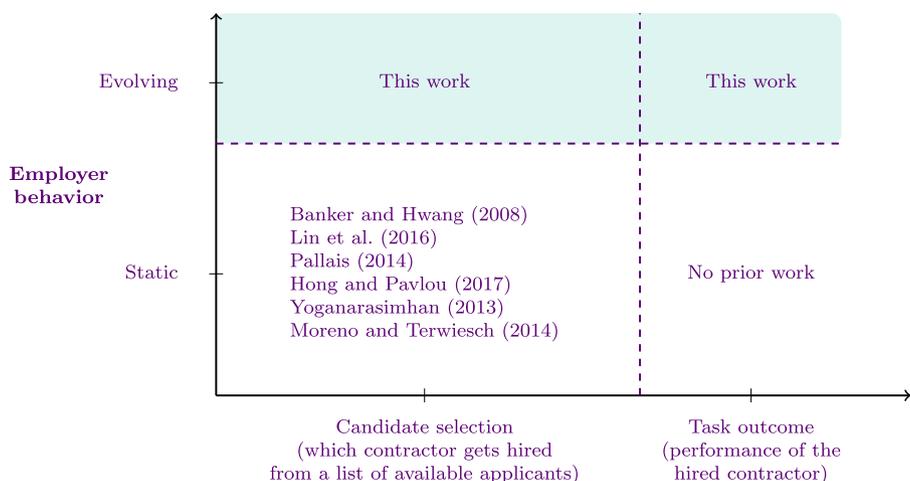
Figure 2 compares this work with the most relevant papers: no prior work investigates how employer hiring preferences evolve and whether they learn to make better hiring choices over time.

2.2. Employer Learning Through Exploration and Exploitation

As employers use an online labor market, their diverse hiring choices result in different observed outcomes. Following a learning curve (Yelle 1979, Schilling et al. 2003), employers can learn from exploration to estimate better the fit between a task’s requirements and an applicant’s abilities. Exploratory learning can happen in two ways: when employers observe a successful outcome, they can reinforce (i.e., use applicant information similarly in subsequent hirings), *exploiting* those successful hiring behaviors (Mom et al. 2007, Lee 2019) that maximized their expected return (Erev and Roth 2014). But when they observe an unsuccessful outcome, they can *explore* by changing how they use information in subsequent hirings in hopes of achieving better outcomes.

In particular, exploitation allows employers to refine existing certainties (Mom et al. 2007) and increase decision reliability (Levinthal and March 1993, Good and

Figure 2. (Color online) Learning from Hiring Decisions in Labor Markets



Notes. This research extends prior research by incorporating two important aspects: (1) Employers and platforms care about outcomes, not just hiring choices. (2) Employers evolve and potentially learn to make better hiring choices.

Michel 2013) by refining what they already know (Holmqvist 2004, Good and Michel 2013), for example, in a new context. On the other hand, unsuccessful hiring choices will lead employers to keep exploring and augment what they already know (Skinner 1990, Gibson and Birkinshaw 2004, Mom et al. 2009, Rogan and Mors 2014, Lee and Meyer-Doyle 2017, Rosing and Zacher 2017, Lee 2019) through a trial-and-error process (Hull 1930, Jueptner et al. 1997, Evans et al. 2000). Because exploration yields uncertain and often negative returns (March 1991, p. 85), and because exploitation usually has “positive, proximate, and predictable” outcomes, we observe that

Proposition 1. *Employers will initially explore the new market environment; over time, they will learn from observed outcomes by adjusting hiring behaviors that lead to unsuccessful outcomes and by reinforcing (exploiting) hiring behaviors that lead to successful outcomes. As a result, the initial exploration choices will be, on average, less successful than the later exploitation ones.*

2.3. Learning to Use Market Information

What attributes do employers use when making hiring choices in online labor markets?

Like new users of other online platforms, new employers in online labor markets confront an assortment of more familiar and less familiar information when making hiring decisions. We focus on two important attributes in hiring choices (Yoganarasimhan 2013, Moreno and Terwiesch 2014, Kokkodis and Ipeirotis 2016)—price and reputation.

As in any market, expense for the employer and revenue for the contractor is crucial. Research in consumer psychology (Shiv et al. 2005) and the neural representations of the brain (Plassmann et al. 2008) shows that consumers associate quality with higher prices. Higher prices signal higher quality in multiple product categories, including frequently purchased convenience goods (Caves and Greene 1996) and relatively cheap products (Rao and Monroe 1989).

In online labor markets, price is a fundamental attribute of a contractor’s bid and is clearly observable (Grewal et al. 1994). Expensive contractors are likely of higher quality and perform better than cheap contractors (Svveaney 1999). As such, pricing has attracted considerable research attention. For example, tasks of higher value (Snir and Hitt 2003) and longer duration (Gefen et al. 2016) attract higher bids but of lower average quality. Sealed bid auctions attract more contractors, but open bid auctions create higher surpluses for the employers (Hong et al. 2015). As bid price dispersion increases, so do employer indecision and contractor regret, hence hurting matching (Zheng et al. 2016). Even further, employers tend to be more price

sensitive in fixed-price contracts than in hourly contracts (Lin et al. 2016).

New online labor market employers can bring considerable experience using price information for transactions in general and perhaps even for hiring in offline settings. Price, however, is noisy and imperfectly correlates with product (or service) quality (Gerstner 1985, Erdem et al. 2008), especially when individuals experience a new context. As a result, some might initially try to *explore* lower-cost options (Mom et al. 2007). (Section 6 discusses why some employers might be more prone to exploring such lower-cost options.) For instance, some new employers may select lower-than-average price contractors to avoid overpaying or to try to beat the market (Grinold and Kahn 2000). Sometimes these lower-price options do result in a good deal for the employer; but, more often, they result in poorer performance (Kokkodis and Ipeirotis 2016). Hence, compared with more experienced employers who have developed increased familiarity with the value of price in this new context and are reinforcing prior successful behavior (Proposition 1), relatively newer employers who are still exploring the value of price in this new context might often make decisions that result in less successful outcomes:

Hypothesis 1. *Employers who explore cheaper contractors will likely have less successful than average outcomes.*

Besides price, online labor markets recognize employers’ need for information about a contractor’s past performance and include some measure of contractor reputation. Current reputation systems allow contractors to receive feedback for the tasks they complete (Filippas et al. 2018). As contractors complete tasks, feedback scores typically accumulate to generate a contractor’s reputation on the platform (Rahman 2018). These contractor reputation scores usually predict future performance (Liu et al. 2008, Danescu-Niculescu-Mizil et al. 2009, Lu et al. 2010, Kokkodis and Ipeirotis 2016, Kokkodis 2021), even in markets with inflated scores (such as online labor markets; Filippas et al. 2018, Rahman 2018). Hence, it is not a surprise that reputation drives the choice of contractors in online labor markets (Yoganarasimhan 2013, Kokkodis 2021). In fact, even having a reputation (compared with being new in the market) significantly improves a contractor’s current (Lin et al. 2016) and subsequent hiring chances (Pallais 2014).

Unlike price, reputation is a platform-specific signal that new employers might initially be less familiar with as these reputation scores do not have exact equivalents outside online labor markets. In fact, although practically all online labor markets include some reputation scores, the details underlying how they work differ idiosyncratically by platform. Hence, some employers might initially explore by discounting the value of reputation and, as mentioned above,

try to hire cheaper contractors who usually have lower (or no) reputation (Moreno and Terwiesch 2014, Pallais 2014). But because contractors with lower reputation scores on average underperform those with higher scores (Kokkodis and Ipeiritos 2016), such lower-reputation exploratory hires will, on average, result in less successful outcomes:

Hypothesis 2. *Employers who explore hiring lower-reputation contractors will likely have less successful than average outcomes.*

Thus, price and reputation signals contain information and the two signals correlate as contractors with higher reputation charge higher premiums (Yoganarasimhan 2013, Moreno and Terwiesch 2014, Pallais 2014). *But how exactly do employers trade off price and reputation over time?*

Consider a group of employers who choose to ignore poor reputation and focus only on lower prices. Stochastically, Hypotheses 1 and 2 suggest that such contractors are more likely to perform poorly. Following Proposition 1, some of those employers whose hired choices underperformed will learn from their bad experience and readjust how they assess market information (e.g., price and reputation). A different group of employers might choose to incorporate reputation even at the cost of slightly higher prices. Such reputable workers are more likely to perform well (Liu et al. 2008, Danescu-Niculescu-Mizil et al. 2009, Lu et al. 2010, Kokkodis and Ipeiritos 2016, Kokkodis 2021). As a result, these employers will be satisfied with their hired choices and, as per Proposition 1, reinforce their hiring behavior. Combined, these dynamics suggest that the first group of employers will move from lower to higher reputation and from cheaper to more expensive contractors, whereas the second group will continue hiring highly reputable contractors who might not be as cheap. Hence, we argue the following:

Hypothesis 3. *Over time, employers learn to increase (or reinforce) their reliance on contractor reputation and decrease (or reinforce) their price sensitivity. This behavior results in more successful than average outcomes.*

3. Modeling Evolving Employer Hiring Choices and Outcomes

Testing these hypotheses requires a framework that models employer evolution and, at the same time, tracks hiring outcomes.

3.1. The Challenges of Simultaneously Modeling Hiring Choices and Outcomes

While important, incorporating both employer evolution and hiring outcomes is difficult—no prior research study includes both (Figure 2). One possible reason for this is that hiring choices and task outcomes occur at

different times and in different populations (i.e., task applicants vs. hired contractors). Employers first choose a contractor to hire, and then they observe the outcome of their collaboration with the chosen contractor. As a result, a model needs to capture first whether a given contractor gets hired (i.e., a binary outcome “hired” versus “no-hired” across task applicants), and second, the performance score of the hired contractor (continuous outcome, across hired contractors).

We propose a two-component framework that directly models these steps (Figure 3). In the Component A, a Hidden Markov Model captures employer evolution across various hiring-behavior states. As a result, this component provides up-to-date estimates of employers’ propensities to hire contractors successfully. Component B uses state-specific conditional logit models to explain how successful and unsuccessful employers choose contractors. With these two components, the framework *compares employers’ hiring behaviors as they evolve*.¹

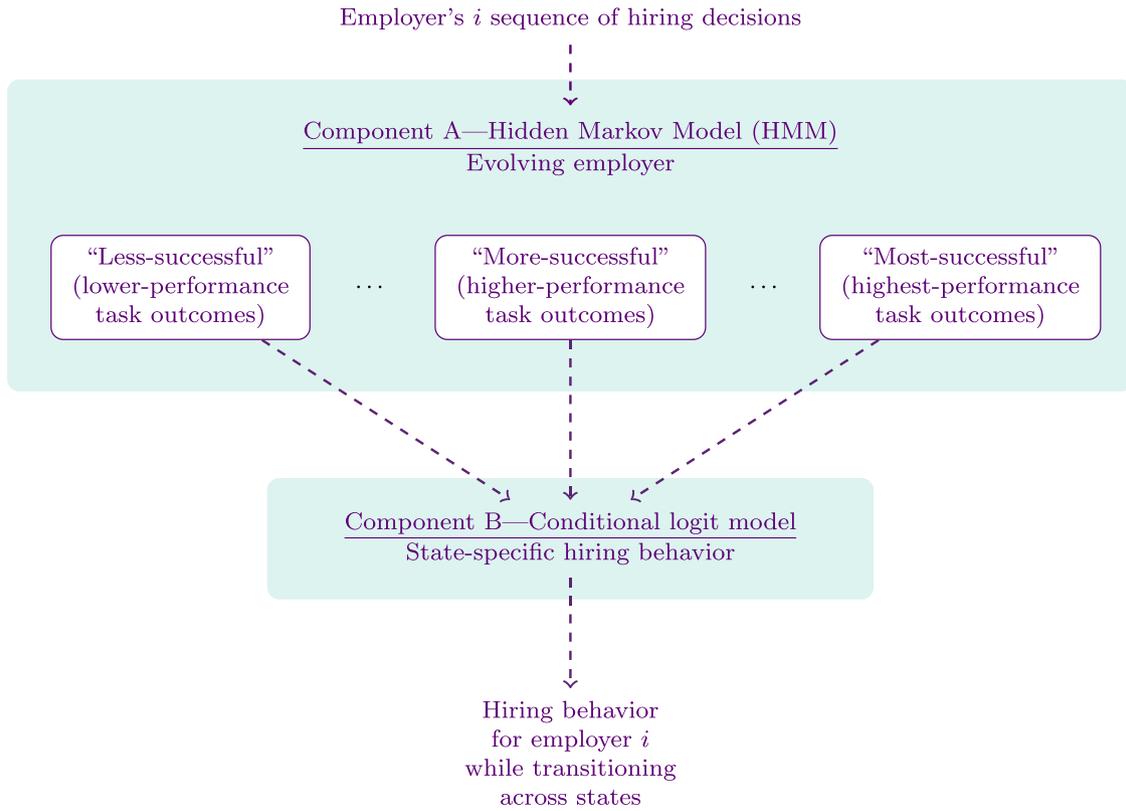
3.2. Component A: Modeling an Employer’s Evolving Hiring Behavior

An online labor market transaction starts with an employer creating a new opening. The opening describes task requirements and contractor characteristics (such as skills and experience) that the employer seeks. Contractors who are searching for opportunities and see a fit between their interests and the task opening can submit an application. For each contractor j that applies for a task o the employer i makes a hiring decision. Once the contractor completes the task, the employer rates the performance of the *hired* contractor (y_{io}).

Each employer who joins the platform has an unobserved process of evaluating available market information to hire contractors. By hiring contractors and observing task outcomes, some employers might learn and adjust their hiring behaviors (Section 2.3). This evolution suggests that employers transition across a series of states, such as “less-successful,” “more-successful,” “most-successful.” Even though we do not observe these states, we observe the outcome of each completed task (i.e., the performance rating y_{io}), which correlates with the underlying employer’s hiring behavior.

An HMM naturally models such an employer evolution across a set of states $\mathcal{S} = \{s_1, \dots, s_K\}$. At any given point in time t , each employer i operates from a state $S_t \in \mathcal{S}$ that represents i ’s propensity to hire successfully. Upon completion of a new task, the employer emits an observation y_{it} (i.e., the performance rating of the hired contractor at time t —we consider that the completion of a new task defines a unit of time; as a result, we omit the task subscript o because the time subscript t subsumes it). Each state models these observations through continuous probability distributions, $Y_{it} \sim f(S_{it}), Y_{it} \in [0, 1], S_{it} \in \mathcal{S}$. (Y_{it} is a random variable that describes observations y_{it} .)

Figure 3. (Color online) A Two-Component Framework of Employer Learning



Notes. Component A models the evolution of employers across various hiring-behavior states that represent different propensities of observing successful outcomes. Component B captures the hiring behavior of each employer within each state. The result allows for a comparison of hiring behaviors as employers transition from a less successful to a more successful state.

3.2.1. Model Structure. A complete definition of an HMM requires (1) a matrix T of the transition probabilities between states, (2) a matrix E of the state-specific emission probability distributions of the observed outcomes, and (3) a vector of initial state probabilities π .

We assume an employer- and state-specific transition matrix T . We allow the history of each employer i up to time $t - 1$ (captured by a vector of covariates \mathbf{Z}_{it-1} discussed in Online Appendix B.2) to affect the employer’s propensity to transition to a new state. These employer-history effects are also state specific: given \mathbf{Z}_{it-1} , the transition probability for employer i to move from a state S_{it-1} to a state S_{it} at time t is

$$\lambda_{\boldsymbol{\gamma}_k, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_k, S_v} = \Pr(S_{it} = S_v | S_{it-1} = S_k; \boldsymbol{\gamma}_k, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}), \quad (1)$$

where

$$\lambda_{\boldsymbol{\gamma}_k, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_k, S_v} = \begin{cases} \Pr(\boldsymbol{\gamma}_k \mathbf{Z}_{it-1} + \varepsilon_i \leq \alpha^1) = \Lambda(\alpha^1 - \boldsymbol{\gamma}_k \mathbf{Z}_{it-1}), & \text{if } v = 1 \\ \Pr(\alpha^{v-1} < \boldsymbol{\gamma}_k \mathbf{Z}_{it-1} + \varepsilon_i \leq \alpha^v) \\ = \Lambda(\alpha^v - \boldsymbol{\gamma}_k \mathbf{Z}_{it-1}) - \Lambda(\alpha^{v-1} - \boldsymbol{\gamma}_k \mathbf{Z}_{it-1}), & \text{if } 1 < v < K \\ \Pr(\boldsymbol{\gamma}_k \mathbf{Z}_{it-1} + \varepsilon_i > \alpha^{K-1}) = 1 - \Lambda(\alpha^{K-1} - \boldsymbol{\gamma}_k \mathbf{Z}_{it-1}) & \text{if } v = K. \end{cases}$$

In Equation (1), $\boldsymbol{\gamma}_k$ is the vector of coefficients of state s_k that define the weights of \mathbf{Z}_{it-1} in estimating the transition probability to state s_v . We estimate this probability through an ordered logit formulation, where ε_i is the unobserved error term that we assume to follow a logistic distribution (i.e., $\varepsilon_i | \mathbf{Z}_{it-1} \sim \text{Logistic}(0, 1)$), Λ is the logit function, and $\boldsymbol{\alpha} = [\alpha^1, \alpha^2, \dots, \alpha^{K-1}]'$ are state-specific thresholds (Wooldridge 2010). As a result, the transition matrix at time t for employer i is

$$T(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}) = \begin{bmatrix} \lambda_{\boldsymbol{\gamma}_1, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_1, S_1} & \lambda_{\boldsymbol{\gamma}_1, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_1, S_2} & \dots & \lambda_{\boldsymbol{\gamma}_1, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_1, S_K} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{\boldsymbol{\gamma}_K, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_K, S_1} & \lambda_{\boldsymbol{\gamma}_K, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_K, S_2} & \dots & \lambda_{\boldsymbol{\gamma}_K, \boldsymbol{\alpha}, \mathbf{Z}_{it-1}}^{S_K, S_K} \end{bmatrix}, \quad (2)$$

where $\boldsymbol{\gamma} = [\boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_K]'$.

The emission matrix E describes the conditional probabilities of observed outcomes given the current state of a user i :

$$\Pr(Y_{it} | S_t = s_k; \boldsymbol{\theta}_k) = f(\boldsymbol{\theta}_k), \quad (3)$$

where $f(\cdot)$ is a truncated continuous probability distribution (e.g., truncated Gaussian) and $\boldsymbol{\theta}_k$ is the parameter vector of this distribution. (Online Appendix B.3

discusses alternative distribution choices for f .) The emission matrix E is then as follows:

$$E(\theta) = [f(\theta_1), f(\theta_2), \dots, f(\theta_K)], \quad (4)$$

where $\theta = [\theta_1, \theta_2, \dots, \theta_K]'$.

Finally, the framework estimates the probability vector $\pi = [\pi_1, \dots, \pi_K]'$ and the parameter vectors γ, α, θ through maximum likelihood estimation. Online Appendix B presents the derivation of the global likelihood and the parameter estimation details.

3.3. Component B: Within-State Employer Hiring Behavior

The utility that an employer i in state $s_k \in \mathcal{S}$ derives by hiring a contractor j for a task o is

$$u_{io}^j(s_k) = \beta_{s_k} X_{jo} + \kappa_o + \varepsilon_{io}^{js_k}, \quad (5)$$

where X_{jo} is the vector of characteristics of the contractor j for task o , κ_o captures static task-specific effects, and $\varepsilon_{io}^{js_k}$ is the unobserved error term. Similar to prior literature on hiring decisions (Moreno and Terwiesch 2014, Chan and Wang 2017), we describe this utility through a conditional logit model (McFadden 1973). The conditional logit model assumes that consumers make choices among a set of products in a way that maximizes their utility functions (Guadagni and Little 1983). In the context of this work, employers choose the candidate that maximizes their state-specific utility function (Equation (5)).

The conditional logit model assumes that the error term $\varepsilon_{io}^{js_k}$ follows a type I extreme value distribution; hence, the probability of choosing contractor j for task o is

$$\Pr(H_{io}^j = 1 | X_{1o}, X_{2o}, \dots, X_{jo}, s_k) = \frac{\exp(\sum_{s_k \in \mathcal{S}} \mathbb{1}_{s_k} \beta_{s_k} X_{jo})}{\sum_{d=1}^J \exp(\beta_{s_k} X_{do})}, \quad (6)$$

where H_{io}^j is an indicator function that is one when contractor j is considered to get hired for task o and $\mathbb{1}_{s_k}$ is an indicator function that is one when the current state of the employer is s_k . The vector of utility weights β_{s_k} identifies the importance of each characteristic in X_{jo} when employer i is at state $s_k \in \mathcal{S}$.

4. Empirical Setup and Results

To empirically test our theoretical framework, we analyze a data set of hiring decisions from a major online labor market. The focal data set includes 238,364 hiring decisions made by 2,932 employers on 10,385 tasks. For all employers, we observe their complete history (i.e., from the time they joined the platform) and how they evolve over consecutive hiring decisions and completed tasks. The data covers diverse tasks from various categories, including administrative, sales, marketing, writing, data science, and web development (Table 15 in Online Appendix E).

In this market, employers and contractors participate globally. Employers begin a new transaction by posting task openings with specific skill requirements and a respective task description. Available contractors then can apply without cost. On average, each task opening in our data set receives 52 applications (median 35, Table 14 in Online Appendix E). Employers then hire from the available applicants. Once the hired contractor completes the task, the employer evaluates the contractor's performance.

The two-component framework (Figure 3) trains on two different data sets. The HMM component trains only on hired contractors, as it requires task-outcome information. We measure task outcomes through privately assigned feedback from employers. (Employer ratings are only visible to the platform.) Online Appendix B describes the details and the results of this process. On the other hand, the conditional logit model trains on all applications. Hence, it allows us to study how employers choose contractors as they evolve across different ability states. The remainder of this section focuses on the results of the conditional logit model.

4.1. Focal and Control Variables

Several measures (vector X_{jo}) may affect employer hiring choices.

4.1.1. Focal Variables. The contractor's profile lists the feedback score at the time of application. This score reflects the accumulated public feedback that prior employers have given the contractor for all previously completed tasks on the platform (*feedback score*). Similarly, contractors reveal their premiums at the time of application. Because online labor markets are highly heterogeneous in terms of tasks and task requirements, bid prices are only comparable within tasks (Kokkodis et al. 2015). Hence, to study how employers respond to price, we create a relative, task-specific measure of the available bids: we consider whether each bid price is above the median *task-specific bid price* (*bid price*).

4.1.2. Task-Varying Variables. To better isolate the effect of the focal variables, we control for the observed history of each contractor on the platform through the number of completed tasks (*contractor completed tasks*) and the number of completed hours of work (*hours of work*; Chan and Wang 2017). Furthermore, we control for the number of tasks that a contractor is concurrently working on at the time of application (*assigned tasks*) and the self-reported years of contractor experience (*contractor experience*).

Besides platform history, we need to control for other contractor characteristics that often affect employer decisions. Specifically, the models control for a

contractor’s gender (*gender*; Chan and Wang 2017) and location, skills, and qualifications (*countries point-wise mutual information* (PMI), *skills inner product* (IP), *certifications PMI*; see Online Appendix E.1). Finally, we also control for timing effects by including the order in which a contractor’s application arrives (*time of application*) and for employer preference effects by annotating contractors that the employer invited to apply (*invited*; Chan and Wang 2017).

Table 1 summarizes the dependent, focal, and control variables. Online Appendix E presents additional data statistics, including a short discussion on feature engineering (Section E.1 in Online Appendix E) and correlations between these variables (Figure 11 in the online appendix). We log-transform variables with long tails. We standardize all nonbinary variables for faster convergence and easier interpretability. All control variables vary across applicants within each task; otherwise, task fixed effects would absorb these.

4.2. Results

Which hiring choices are more likely to yield successful outcomes, and how do employers adjust their hiring behaviors to make such choices?

Component B’s (Figure 3) conditional logit model estimates each employer’s behavior within each state. The main specification (Equation (6)) controls for observed and unobserved confounding factors that could endogenize the estimates of the focal variables. Specifically, task-varying variables control for the observed heterogeneity across the set of available candidates for each task (Table 1). Opening fixed effects accounts for private information that employers have when making hiring decisions; hence, they eliminate the time-invariant portion of the unobserved error. Opening fixed effects further eliminates any employer-specific, category-specific, and

time-specific static unobserved effects as these quantities do not vary within tasks.

To identify the parameters of Equation (6), we first need to estimate each employer’s most likely state at the time of the hiring decision (i.e., $s_k \in \mathcal{S}$). Online Appendix B details the HMM implementation and estimation process. Through this analysis, the three latent states that best describe the employers’ hiring behavior in our data (Figure 7 in Online Appendix B) are as follows:

- Less-successful state: employers’ decisions result in unsuccessful task outcomes. (Employers hire contractors that, on average, perform 0.53 out of 1.)
- More-successful state: employers’ decisions result in more successful task outcomes. (Employers hire contractors that, on average, perform 0.82 out of 1.)
- Most-successful state: employers’ decisions result in successful task outcomes. (Employers hire contractors that, on average, perform 0.99 out of 1.)

For these estimated latent states $\mathcal{S} = \{ \text{Less-successful, More-successful, Most-successful} \}$, we identify the coefficients of Equation (6). Table 2 shows the results for each of the state-specific focal variables. (Table 16 in Online Appendix F shows the coefficients of the complete vector $\beta_{s_k} \forall s_k \in \mathcal{S}$.) Column (A1) in Table 2 shows the average effect of the two focal variables across states. Column (A2) splits this effect into state-specific effects. First, the decrease in Akaike information criterion score from column (A1) to column (A2) suggests that the state-specific estimates better explain the observed variance in hiring choices—a log-likelihood ratio test (Wooldridge 2010) rejects the null hypothesis that the static model (column (A1)) fits the data better ($p < 0.001$), providing further statistical evidence that modeling employer evolution through an HMM better explains the observed variance in hiring choices.

Table 1. Descriptive Statistics of the Dependent, Focal, and Control Variables

Variable	Units	Mean	Median	StD	Min	Max
Dependent variable						
<i>Hired</i> (H_{io}^i)	Hired = 1	0.04	0	0.2	0	1
Focal variables						
<i>Feedback score</i>	Index	0.94	0.98	0.1	0	1
<i>Bid price</i>	Above median = 1	0.37	0	0.48	0	1
Task-varying variables						
<i>Assigned tasks</i>	Count	3.4	2	5	0	90
<i>Skills IP</i>	Count	1	1	1.2	0	16
<i>Certifications PMI</i>	Index	2.4	2.3	1.7	−0.71	7
<i>Hours of work</i>	Hours	551	44	1,401	0	27,662
<i>Time of application</i>	Ordinal	25	15	31	0	263
<i>Gender</i>	Male = 1	0.81	1	0.39	0	1
<i>Contractor completed tasks</i>	Count	5	0	14	0	356
<i>Contractor experience</i>	Years	4.5	4	4.1	0	30
<i>Invited</i>	Invited = 1	0.12	0	0.33	0	1
<i>Countries PMI</i>	Index	−0.5	−0.46	0.43	−4.9	5

Note. StD, standard deviation (see Online Appendix E.1 for additional details).

Table 2. State-Specific Effects of Feedback Score and Bid Price on Hiring Choices

DV: Hire	All employers		Evolving employers (with HMM) (A3)
	(without HMM) (A1)	(with HMM) (A2)	
Bid price	-0.131*** (0.023)		
Feedback score	0.288*** (0.018)		
Bid price × less-successful		-0.350*** (0.053)	-0.374*** (0.058)
Bid price × more-successful		-0.146** (0.048)	-0.173*** (0.051)
Bid price × most-successful		-0.053. (0.030)	-0.038 (0.036)
Feedback score × less-successful		0.142*** (0.033)	0.142*** (0.037)
Feedback score × more-successful		0.248*** (0.037)	0.266*** (0.040)
Feedback score × most-successful		0.379*** (0.027)	0.391*** (0.033)
Task-varying variables	Yes		
Task-varying variables × less-successful		Yes	Yes
Task-varying variables × more-successful		Yes	Yes
Task-varying variables × most-successful		Yes	Yes
Opening FE	Yes	Yes	Yes
Observations	238,364	238,364	174,873
AIC	52,833	52,763	39,113

Notes. Results of conditional logit models. The first two columns show results of the complete data set (all employers). The last column focuses only on employers who transition across two or more states. Clustered standard errors in parentheses. Constant term is estimated but omitted from the table. Tables 16 and 17 in the online appendix show the state-specific coefficients for all task-varying variables. AIC, Akaike information criterion; DV, Dependent Variable.

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; . $p < 0.1$.

Second, as employers move from less-successful to more-successful states, they increase their reliance on the observed feedback score and tend to hire relatively more expensive contractors (Hypothesis 3). The focal coefficients differ both practically and statistically ($p < 0.001$) across states. Specifically, in line with Hypothesis 1, the coefficient of bid price is significantly smaller for employers in more successful states: an employer in the most-successful state is 85% less price sensitive than an employer in the less-successful state. In addition, in line with Hypothesis 2, the coefficient of feedback score is significantly larger for employers in more-successful states: an employer in the most-successful state values reputation 267% more than an employer in the less-successful one.

The results in column (A2) include all employers, even those who never transition to different states. As a result, they might represent disjoint sets of employers with different hiring preferences and not employers who evolve across states. To investigate, we consider only *evolving employers* who transition across two or more states. This subsample compares how *the same employers* hire in less-successful versus more-successful versus most-successful states. Hence, it identifies whether the changes in employers' hiring behaviors drive the observed results. The results (column (A3) of Table 2) are

consistent with the entire sample results; because they focus on evolving employers, they provide stronger support for Hypothesis 3: as employers move from less successful to more successful hiring choices, they learn to adjust their hiring behaviors by relying more on feedback score and becoming less price sensitive.

4.3. Additional Evidence in Support of the Theoretical Framework

The previous analysis supports Hypotheses 1, 2, and 3. To strengthen this empirical support, we investigate whether data evidence also aligns with the theoretical arguments that derived these hypotheses (Section 2.3). Specifically, our conceptualization argues that (1) lower prices and lower reputation scores yield, on average, less successful results and (2) employers learn and adjust their hiring behavior (by exploring) after unsuccessful outcomes, whereas they reinforce their hiring behaviors (by exploiting) after successful ones. Next, we provide empirical evidence in support of these arguments.

4.3.1. Further Support for Hypotheses 1 and 2. Are employers who choose cheaper contractors with a lower reputation more likely to yield unsuccessful outcomes?

The main analysis in Table 2 shows that employers who are more likely to make successful hiring choices choose, on average, more expensive contractors with a higher reputation (supporting Hypotheses 1 and 2). An alternative way to show that indeed such choices correlate with more successful outcomes is to estimate the average bid price and feedback score for various outcomes. First, we separate successful and unsuccessful outcomes by adapting the internal annotation of the focal platform: if the employer privately rates the performance of the contractor with a score greater or equal to 80%, the platform labels that outcome as successful; otherwise, the platform labels the outcome as unsuccessful. For simplicity, we call successful outcomes “hire-positive” and unsuccessful ones “hire-negative.” (Figure 1(a) shows how this threshold separates hire-positive from hire-negative outcomes.) Figure 4, (a1) and (a2) show that successful outcomes correlate with a significantly higher feedback score ($p < 0.01$) and bid price choices ($p < 0.05$), hence providing additional support for Hypotheses 1 and 2.

4.3.2. Further Support for Hypothesis 3. Do employers learn to adjust their hiring preferences?

Column (A3) in Table 2 shows that, indeed, many employers transition across states; when they move from less- to more-successful states, they adjust their hiring criteria in accordance with Hypothesis 3. An alternative way to test our arguments is to track how

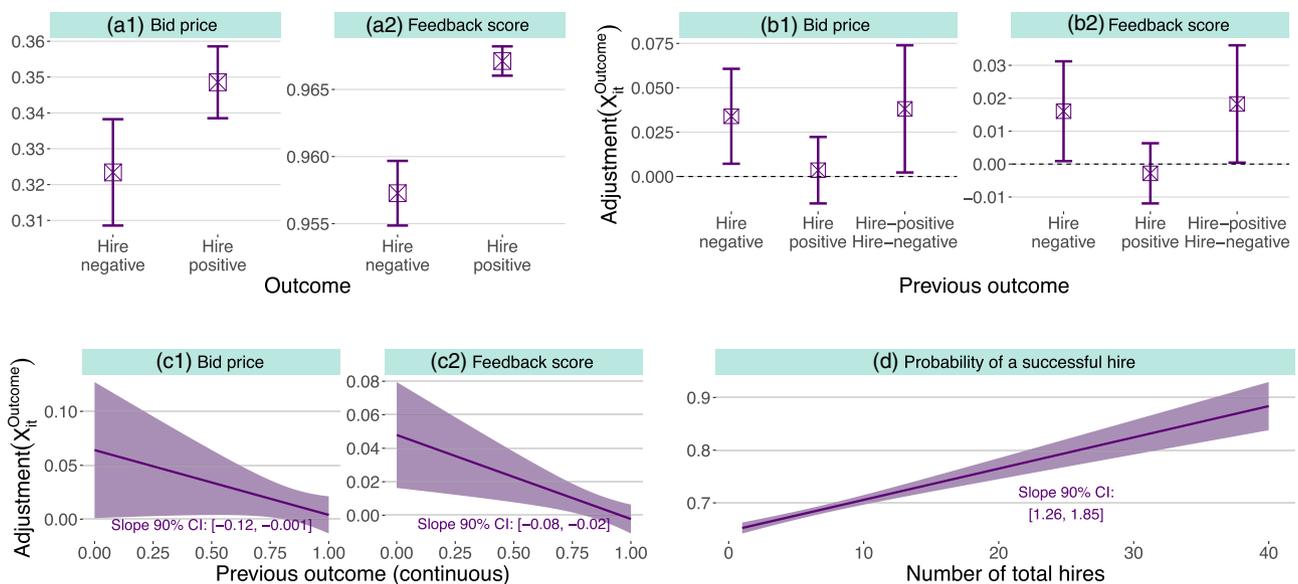
employers react after successful and unsuccessful hiring decisions. Given an employer i and a hiring outcome at time t , we define employer i 's hiring adjustment of feature X as follows:

$$\text{Adjustment}(X_{it}^{\text{Outcome}}) = X_{it} - X_{it-1}^{\text{Outcome}}. \quad (7)$$

We estimate the average adjustment for $X \in \{\text{bid price, feedback score}\}$ and $\text{Outcome} \in \{\text{hire-positive, hire-negative}\}$. Figure 4, (b1) and (b2) show that, indeed, after unsuccessful outcomes, employers explore and hire contractors with significantly ($p < 0.05$) better feedback scores who are more expensive. On the other hand, after successful outcomes, employers exploit and reinforce their behavior (insignificant adjustment, $p > 0.1$). The figures also show that a similar adjustment occurs when employers observe a less successful (hire-negative) outcome after previously observing a more successful (hire-positive) one. Combined, this evidence provides empirical support that, on average, employers indeed learn and adjust their behaviors in accordance to Hypothesis 3.

Figure 4, (c1) and (c2) show the linear relationship between outcome adjustment (Equation (7)) and the continuous version of each employer's previous hiring outcome. The figures show that for both bid price and feedback score, the slope is negative and significant (at least $p < 0.1$), showing that larger adjustments correlate with a worse performance of the previous hire. As a result, these results align with those of Figure 4, (b1)

Figure 4. (Color online) Additional Support for Hypotheses 1, 2, and 3



Notes. Higher bid price and feedback score correlate with more successful outcomes (a1), (a2). After unsuccessful outcomes, employers explore by choosing contractors with higher bid price and feedback score (b1), (b2); the y -axis shows the adjustment between consecutive hires of the same employer). The more negative the outcome, the larger is the adjustment (c1), (c2). Overall, the more hiring choices an employer makes, the higher is the likelihood of a successful outcome (d). Error bars show 95% bootstrapped confidence intervals. CI, confidence interval.

and (b2); they further support our core arguments about employer learning.

Finally, Figure 4(d) shows employer learning through a simpler regression: the probability of a successful outcome increases ($p < 0.05$) as employers make more hiring choices. This increase suggests that over time employers move from less to more successful hiring choices, further corroborating our findings in support of Hypothesis 3.

In summary, our analyses in this section support our conceptualization and show that many employers do indeed learn to adjust their preferences and hire more successfully over time. On the other hand, support for the proposed mechanisms does not necessarily exclude alternative mechanisms that could also be driving the results. Section 5 and Online Appendix A test 10 potential alternative mechanisms.

5. Alternative Explanations

To increase confidence in the hypothesized conceptualization, we

- Incorporate the choice of an employer to participate or exit the market (using a customized exit-state HMM and Heckman two-stage selection models, Online Appendix A.1);
- Focus only on repeat employers (using subsample analyses, Online Appendix A.2);
- Control for employer hiring ability (using instrumental variables, Online Appendix A.3);
- Examine whether more experienced or better employers attract better contractors and as a result learning does not actually occur (normalization of reputation within each opening, distribution of applicants, Online Appendix A.4);
- Consider alternative measures of outcome success (such as rehires, Online Appendix A.5);
- Control for task characteristics (using matched samples in conditional logit models, Online Appendix A.6);
- Explore how hiring abilities differ between technical and nontechnical tasks (using subsample analyses, Online Appendix A.7);
- Investigate how focused or diverse employers drive the results (using subsample analyses, Online Appendix A.8);
- Model employer evolution and hiring choices with alternative approaches (using clustering algorithms and various fixed effects specifications, Online Appendix A.9); and
- Consider different specifications that use employer fixed effects and control for time-varying employer attributes (e.g., experience) that opening fixed effects absorb in our main specification (Online Appendix A.10).

Table 3 summarizes these tests, which corroborate our main findings and provide additional evidence

that employer learning occurs as, on average, employers transition from choosing cheaper contractors with a lower reputation to choosing more expensive contractors with a higher reputation.

6. Potential Mechanism: Employer Confidence

The empirical analysis clarifies the dynamics of learning in online labor markets. However, our conceptualization in Section 2 does not explain why some platform participants might be more likely to explore lower-cost options.

One possible explanation is confidence (i.e., belief in success; Snyder and Lopez 2009). An individual's confidence structures their attitude toward information (Kreye et al. 2012) as it governs the perception of loss associated with specific actions (Forlani and Mullins 2000). Higher confidence lowers perceptions of risk (Busenitz and Barney 1994, Palich and Bagby 1995, Robinson and Marino 2015) and results in "optimistic judgments of risk" (Kahneman and Lovallo 1993, p. 29).

6.1. Confidence in Online Labor Markets

In an online labor market, an employer's confidence and the resulting risk assessment may influence how an employer assesses available market information. Employers with higher confidence will be more likely than employers with lower confidence to risk and try to beat the market (Grinold and Kahn 2000) by looking for cheap contractors (with likely lower reputation scores—Section 2.3). But our theory and empirical analysis finds that cheaper contractors on average yield worse outcomes (Section 2.3; Hypotheses 1 and 2; and Figure 4, (a1) and (a2)). As a result, some of these higher-confidence employers who hire unsuccessfully will readjust their confidence and risk attitudes and, hence, learn to better assess the available information sources (Proposition 1, Hypothesis 3).

Conversely, employers with lower confidence may minimize risk by hiring contractors with a higher reputation albeit at a higher price. Their choices will likely yield, on average, more successful outcomes (Section 2.3; Hypotheses 1 and 2; and Figure 4, (a1) and (a2)), which in turn might reinforce their confidence and risk attitudes (exploitation strategies; March 1991, Denrell and March 2001, Menkhoff et al. 2006, Goodman-Delahunty et al. 2010, Lee and Meyer-Doyle 2017). Some of these contractors might then explore more risky choices until they learn to appropriately assess available information (Hypothesis 3; Figure 4, (b1) and (b2)).

These dynamics suggest that over a period of time, the market will observe employers with higher confidence to move from higher-risk (i.e., cheaper contractors with a lower reputation) to lower-risk (i.e., more

Table 3. Investigation of Alternative Explanations

Analysis	Bid price (support for H1)	Feedback score (support for H2)	Discussion
Alternative explanation: self-selection of employers to keep participating drives the results			
Solution: modeling the choice to participate:			
Exit state HMM	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.1, Table 4
Heckman two-stage selection model	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
Alternative explanation: repeat employers always had the observed behavior (i.e., no learning over time)			
Solution: focus only on repeat employers who hire:			
Two or more times	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.2, Table 5
Three or more times	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
Four or more times	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
Five or more times	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
Alternative explanation: the ability of employers who end up staying on the platform drives the observed results			
Solution: instrumental variables capture employer state:			
State of every other employer at hire o	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.3, Table 6
% of employers who survive at hire o	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
% of employers who survive at hire o after outcome	✓ ($p < 0.001$)	✓ ($p < 0.001$)	
Alternative explanation: more experienced employers attract better contractors			
Solution: normalize reputation within openings	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.4, Table 8, Figure 6
Distributions of applicants across employers	L	L	
Alternative explanation: the results are an artifact of our measurement of success			
Solution: rehires as measure of success	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.5, Table 7
Alternative explanation: systematically different types of tasks end up being in different states			
Solution: propensity score matching on task characteristics	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.6, Table 9
Alternative explanation: employers might have different abilities across different types of tasks			
Solution: separate analysis of technical and nontechnical tasks	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.7, Table 10
Alternative explanation: focused or diverse employers drive the results			
Solution: separate analysis of focused and diverse employers	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.8, Table 10
Alternative explanation: results are an artifact of the HMM modeling			
Solution: alternative modeling of Component A (K -means)	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Online Appendix A.9, Table 11
Alternative explanation: results are an artifact of the conditional logit model approach			
Solution: more experienced employers attract better contractors / alternative modeling of Component B:			
Linear models with employer FE	✓ ($p < 0.001$)	✓ ($p < 0.001$)	Section A.10, Table 12 in the online appendix
Logit models with employer FE	✓ ($p < 0.001$)	✓ ($p < 0.001$)	

Note. ✓, Hypothesis supported; (p), the level of significance that the respective coefficient of the most-successful state is different than the respective coefficient of the less-successful state; FE, Fixed Effects; L , suggests support of our core employer learning argument; H1, Hypothesis 1; H2, Hypothesis 2.

expensive contractors with a higher reputation) choices, whereas employers with lower confidence to reinforce their behavior or explore higher-risk choices that, when

unsuccessful, will push them back to lower-risk choices. As a result, the net effect of these dynamics will be a transition from lower to higher contractor reputation

and premium and, on average, from less to more successful hiring outcomes—the invisible hand at work.

6.2. Empirical Evidence of the Confidence Mechanism

The market does not directly observe employers’ confidence levels. However, text analysis can detect concepts related to confidence. Specifically, we use the Linguistic Inquiry and Word Count (LIWC) package to analyze the unstructured text of employers’ task descriptions (LIWC 2018). LIWC measures dimensions such as social processes (Sridhar and Srinivasan 2012), emotion (Hong et al. 2016), popularity (Goes et al. 2014), or anxiousness (Yin et al. 2014) using validated dictionaries. Like other research (Smith-Keiling and Hyun 2019), we measure confidence using the LIWC *clout* metric, defined as the relative social status, confidence, or leadership that people display through their writing or talking (LIWC 2018).

Alternatively, the task description’s *emotional tone* may also correlate with confidence. A higher emotional tone conveys a more positive attitude, whereas a lower emotional tone reveals greater anxiety (LIWC 2018). Given that higher confidence associates with a more positive attitude (Snyder and Lopez 2009), emotional tone might serve as an additional signal of employers’ confidence.

6.2.1. Do Higher-Confidence Hiring Choices Result in Less Successful Outcomes? Figure 5(a) suggests that the answer is probably yes: higher clout ($p < 0.05$) and emotional tone ($p < 0.1$) positively correlate with the likelihood that the hiring outcome will be unsuccessful (Figure 5(a)). This observation aligns with the argument that higher-confidence employers make higher-

risk choices of cheaper contractors with lower reputation scores who, on average, result in less successful outcomes (Section 2.3; Hypotheses 1 and 2; and Figure 4, (a1) and (a2)).

6.2.2. Do Employers Adjust Their Confidence and Risk Attitudes Over Time? To investigate, we compute the employer-specific confidence adjustment at time t as follows:

$$\text{Normalized } C_{it} = C_{it} - C_{i1}, \quad (8)$$

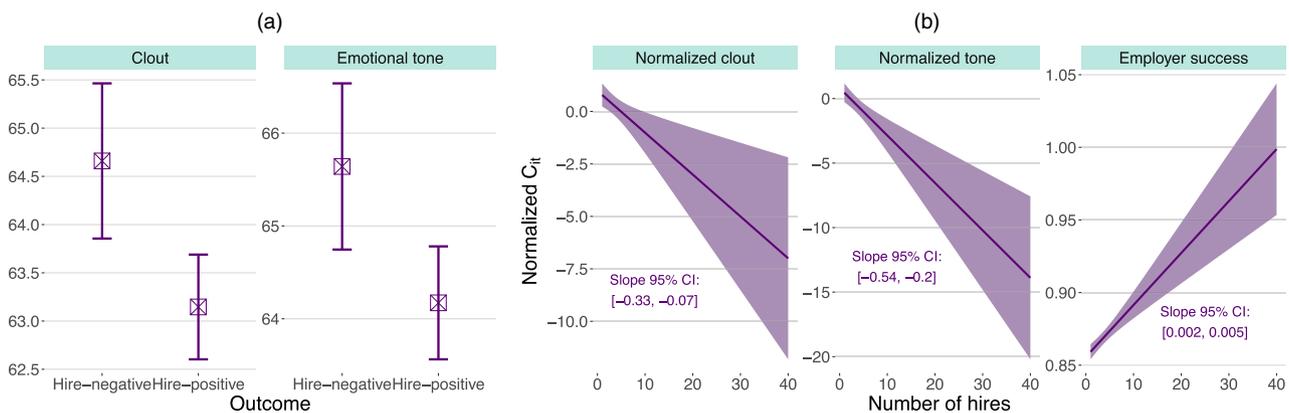
where i is the focal employer and $C \in \{\text{Clout, Emotional tone}\}$. This normalized score proxies how employers adjust their confidence level over time compared with their own initial confidence (comparison level).

As employers hire more contractors, they generally adjust their confidence downward (negative relationship, $p < 0.001$, in the first two panels in Figure 5(b)). Because employers mostly adjust their behaviors after unsuccessful outcomes (Figure 4, (b1), (b2), (c1), and (c2)), the observed effect likely stems from employers of higher-confidence and higher-risk attitudes. Furthermore, confidence readjustment suggests that the transitions from higher- to lower-risk choices should result (on average) in more successful outcomes. Indeed, over time, employers hire contractors who perform better ($p < 0.001$, last panel in Figure 5). (Online Appendix B.4 provides additional empirical evidence in support of the confidence explanation.)

7. Discussion

How do employers adjust their behaviors as they learn to make successful hiring decisions in online labor markets? We theorized that employers initially explore

Figure 5. (Color online) Over Time, Higher-Confidence (Proxied by Clout and Emotional Tone) Employers Adjust Their Hiring Preferences



Notes. Higher-confidence employers (as measured by higher clout and emotional tone) have on average less successful outcomes (a). Over consecutive hiring choices (x - axis), employers adjust their confidence (as measured by clout and emotional tone), and this adjustment correlates with more successful hiring choices (b). CI, confidence interval.

in their assessment of available market information; some—potentially because of higher confidence—will try to beat the market and hire contractors who are cheaper and have a lower reputation. But these contractors are more likely to have poor performance. Hence, employers who observe such unsuccessful outcomes will learn and adjust their hiring behavior by further exploring the value of available market information. On the other hand, employers who hire successfully will reinforce (exploit) their hiring strategies. We empirically test these mechanisms by building a two-component framework (HMM, conditional logit models) that links hiring behaviors with task outcomes. Hiring decisions from a large online labor market provide empirical evidence in support of this theory: employers learn and adjust their hiring behaviors by transitioning from hiring cheaper contractors with a lower reputation to hiring more expensive contractors with a higher reputation. The latter yields significantly more successful results.

7.1. Contributions to Research

This research extends current literature in online labor markets and hiring decisions by investigating a new question (Figure 2): which hiring choices are more likely to yield successful outcomes, and how do employers learn to adjust their hiring behaviors to achieve better outcomes? Through a new theoretical framework, our work argues that, over time, many employers learn to make successful hiring choices. This theoretical framework deepens our understanding of how employers' hiring preferences evolve and guides the empirical identification of the studied effects. Through a comprehensive empirical analysis, our study is the first to show that as employers transition from unsuccessful to successful hiring decisions, they rely more on contractors' feedback scores and become less price sensitive.

This paper is the first to link task outcomes with employer hiring behaviors and show that successful outcomes often associate with higher contractor feedback scores and higher premiums. The two-component framework provides a guideline for both studying hiring decisions in a dynamic context and linking task outcomes with hiring choices. As a result, future research can study matching in digital workplaces through evolving employer and contractor interactions while incorporating task outcomes.

Additionally, this methodological contribution (i.e., the two-component framework) spans beyond the specific context of online labor markets. For instance, the proposed framework can investigate employer-evolving behavior in the offline setting (e.g., through platforms such as LinkedIn or BurningGlass). Task outcomes, in that case, could be promotions and other aspects of career trajectories. Overall, organizations

stand to benefit by applying the two-component framework to study hiring behaviors and outcomes formally.

7.2. Implications for Platform Managers

Given the high potential for online work growth in the coming years (Agile-1 2016, Sundararajan 2016), understanding how employers can improve their likelihood of hiring successfully can provide platform managers with an additional tool to facilitate employer success. In particular, our work provides platforms with an explanation of why some hiring choices are more likely to be unsuccessful. Understanding what makes a successful hiring decision and how employers learn to make successful candidate choices is fundamentally important for an online labor market—both (1) employers who hire successfully and (2) contractors who receive positive feedback are more likely to keep participating in the platform (Jerath et al. 2011, Tripp and Grégoire 2011).

Furthermore, our study provides online labor markets with a data-driven approach (HMM framework) that models the evolving hiring abilities of the market's clients (employers). Through the proposed framework, managers can better understand the market's employer population: at any point in time, managers can identify employers who are struggling to make successful hiring choices. Managers can then intervene and provide these employers with information and guidance on how to minimize uncertainty and hire candidates that are likely to perform well. Even further, and depending on the market's policies, managers can experiment with subsidizing such employers (e.g., by reducing the platform's commission) to further increase their chances of success and subsequent use of the platform.

7.3. Limitations and Future Research

Hiring decisions in online labor markets depend on factors other than price and reputation. For example, employers base hiring decisions on contractor demographic attributes, as they prefer contractors from their own country (Gefen and Carmel 2008, Ghani et al. 2014, Lin and Viswanathan 2015) because communication is easier and trust and familiarity increases when interacting with local contractors (Arora and Forman 2007, Kossinets and Watts 2009). Furthermore, contractor gender (Chan and Wang 2017), skills, and certifications also affect employer choices (Goes and Lin 2013, Kokkodis et al. 2015). In this work, we focus on price and reputation because they can sufficiently illustrate how users learn in online labor markets. Our empirical analysis, however, uses all these additional characteristics as control variables (Table 1); in Appendix F and Table 16 in the online appendix, we briefly mention how their estimated coefficients fit within the context

of this framework. Future research can use this framework to further theorize and explore how employers adjust their hiring behavior across these additional characteristics.

A limitation of our study is that it relies on observational data to investigate how employers learn. The empirical evidence in Section 4 and Online Appendix A supports that at least some employers become successful through the learning path discussed in Section 2.3. Yet, we cannot claim that learning is the only causal path that drives these results. Hypothetically, a field experiment could randomly assign contractors (by drawing from a predetermined pool of contractors with certain characteristics) of varying reputation and bid prices to be chosen by employers. This way, the experiment would enforce choice sets that are almost identical (same number of applicants with similar characteristics) across employers and hence eliminate any employer ability to attract better (or worse) pools of applicants. Next, the platform would observe how employers react after different types of outcomes and whether they adjust their behaviors after observing hiring outcomes. By ensuring that no employer exits the market, the experiment could show whether employer learning is indeed the sole causal path that drives these results.

The above hypothetical experiment has practical and ethical limitations (e.g., random assignment of applicants). Our study, even though observational, takes significant steps to investigate whether employers attract systematically better (or worse) pools of applicants (Appendices A.3 and A.4 and Figure 6 in the online appendix), and whether employers who exit the market affect the results (Online Appendices C, A.1, A.2, and A.3). However, we acknowledge that our observational analysis cannot fully eliminate these mechanisms, which could also explain part of the observed results.

Section 6 suggests that as employers explore market information and encounter unsuccessful outcomes, they adjust their behavior by reducing risk and achieve more successful outcomes. However, some users who take on additional risk are successful; future research could help better identify conditions in which seemingly risky selections of lower-reputation contractors, for example, may be worthwhile.

Additionally, we expect that new users similarly adapt in other platforms; future research can help better understand how users learn and adjust their decisions in such different contexts. For example, platform designers could make prior outcomes more or less salient to users or platform managers could clarify or obscure how user decisions deviate from other users' decisions. Importantly, user adaptation likely depends on many context attributes, such as frequency of use, subjectivity of outcomes, process steps, composition and adaptation of other users, and consequence of unsuccessful outcomes. Future research can clarify these issues and many others.

7.4. Conclusion

Business is increasingly digital with online platforms underlying an ever-growing portion of modern commerce. These digital platforms overcome many market inefficiencies by facilitating transactions between consumers and producers throughout the world. But as new digital markets and platforms emerge, new platform users come in with diverse backgrounds that lead to varied ability to navigate the new markets. As a result, many online platform users struggle to adapt their offline experiences to digital platforms. These struggles are costly for both the user and the platform.

Yet, despite the importance of learning on digital platforms, we know little about how new users adjust their behavior based on experiences on the platform. Our work on employer learning in online labor markets opens a discussion on how platform users can learn and adjust their behaviors as they gain experience. Beyond hiring choices, in digital platforms, such as ride-matching² or hotels,³ many users might need to learn new managerial skills. Our framework provides a guideline on platform learning that other platforms can customize to understand how their users adjust their behavior over time in order to succeed.

Acknowledgments

The authors would like to thank Panos Adamopoulos, Vilma Todri, and Kostas Pelechris for their helpful comments and suggestions.

Endnotes

¹ The two-step process in our context has both similarities and differences with traditional two-stage selection models (Heckman 1979). Like these models, our framework models two stochastic events: the probability of an employer having a given level of hiring ability and the employer's subsequent probability of choosing an applicant (conditional logit). Yet, traditional selection models cannot directly model the first step of our approach because this step does not rely on an observed outcome (e.g., the choice to participate), but instead it estimates *unobserved* ability states. On the other hand, two-stage selection models can capture the choice to participate or exit the market; we implement such models to test alternative explanations in Section A.1.2 in the online appendix.

² See <https://www.uber.com/newsroom/first-trip-tips-riders/>.

³ See <https://money.usnews.com/money/blogs/my-money/2015/06/23/6-things-to-know-when-using-airbnb-for-the-first-time>.

References

- Agile-1 (2016) Gig economy. Accessed November 28, 2020, http://www.hrotoday.com/wp-content/uploads/2016/07/Whitepaper_Agile2016-single.pdf.
- Agrawal A, Hortton JJ, Lacetera N, Lyons E (2015) Digitization and the contract labor market. Goldfarb A, Greenstein SM, Tucker CE, eds. *Economic Analysis of the Digital Economy* (University of Chicago Press, Chicago), 219–250.
- Allon G, Bassamboo A, Çil EB (2012) Large-scale service marketplaces: The role of the moderating firm. *Management Sci.* 58(10): 1854–1872.

- Arora A, Forman C (2007) Proximity and information technology outsourcing: How local are IT services markets? *J. Management Inform. Systems* 24(2):73–102.
- Barr PS, Stimpert JL, Huff AS (1992) Cognitive change, strategic action, and organizational renewal. *Strategic Management J.* 13(S1):15–36.
- Brier E, Pearson R (2018) Upwork's SVP of marketing explains what it takes to perfect an offering that relies on people. Accessed November 28, 2020, <https://techdayhq.com/community/articles/upwork-s-svp-of-marketing-explains-what-it-takes-to-perfect-an-offering-that-relies-on-people>.
- Broyden CG (1970) The convergence of a class of double-rank minimization algorithms. 1. General considerations. *J. Appl. Math.* 6(1):76–90.
- Busenitz LW, Barney JB (1994) Biases and heuristics in strategic decision making: Differences between entrepreneurs and managers in large organizations. *Acad. Management Proc.* 1994(1):85–89.
- Busenitz LW, Barney JB (1997) Differences between entrepreneurs and managers in large organizations: Biases and heuristics in strategic decision-making. *J. Bus. Venturing* 12(1):9–30.
- Caves RE, Greene DP (1996) Brands' quality levels, prices, and advertising outlays: Empirical evidence on signals and information costs. *Internat. J. Indust. Organ.* 14(1):29–52.
- Chan J, Wang J (2017) Hiring preferences in online labor markets: Evidence of a female hiring bias. *Management Sci.* 64(7):2973–3468.
- Chen DL, Horton JJ (2016) Are online labor markets spot markets for tasks? A field experiment on the behavioral response to wage cuts. *Inform. Systems Res.* 27(2):403–423.
- Danescu-Niculescu-Mizil C, Kossinets G, Kleinberg J, Lee L (2009) How opinions are received by online communities: A case study on Amazon.com helpfulness votes. *Internat. Conf. World Wide Web* (ACM, New York), 141–150.
- Denrell J, March JG (2001) Adaptation as information restriction: The hot stove effect. *Organ. Sci.* 12(5):523–538.
- Dequech D (1999) Expectations and confidence under uncertainty. *J. Post Keynesian Econom.* 21(3):415–430.
- Dey D, Fan M, Zhang C (2010) Design and analysis of contracts for software outsourcing. *Inform. Systems Res.* 21(1):93–114.
- Erdem T, Keane MP, Sun B (2008) A dynamic model of brand choice when price and advertising signal product quality. *Marketing Sci.* 27(6):1111–1125.
- Erev I, Roth AE (2014) Maximization, learning, and economic behavior. *Proc. National Acad. Sci. USA* 111(supplement_3):10818–10825.
- Evans JJ, Wilson BA, Schuri U, Andrade J, Baddeley A, Bruna O, Canavan T (2000) A comparison of “errorless” and “trial-and-error” learning methods for teaching individuals with acquired memory deficits. *Neuropsychological Rehabilitation* 10(1):67–101.
- Filippas A, Horton JJ, Golden J (2018) Reputation inflation. *Conf. Econom. Comput.* (ACM, New York), 483–484.
- Forlani D, Mullins JW (2000) Perceived risks and choices in entrepreneurs' new venture decisions. *J. Bus. Venturing* 15(4):305–322.
- Gefen D, Carmel E (2008) Is the world really flat? A look at offshoring at an online programming marketplace. *MIS Quart.* 32(2):367–384.
- Gefen D, Gefen G, Carmel E (2016) How project description length and expected duration affect bidding and project success in crowdsourcing software development. *J. Systems Software* 116:75–84.
- Gerstner E (1985) Do higher prices signal higher quality? *J. Marketing Res.* 22(2):209–215.
- Geva T, Saar-Tsechansky M (2016) Who's a good decision maker? Data-driven expert worker ranking under unobservable quality. *Internat. Conf. Inform. Systems* (Association for Information Systems, Atlanta), 1–16.
- Ghani E, Kerr WR, Stanton C (2014) Diasporas and outsourcing: Evidence from oDesk and India. *Management Sci.* 60(7):1677–1697.
- Gibson CB, Birkinshaw J (2004) The antecedents, consequences, and mediating role of organizational ambidexterity. *Acad. Management J.* 47(2):209–226.
- Goes P, Lin M (2013) Does information really “unravel” understanding factors that motivate sellers to seek third-party certifications in an online labor market. Working paper, Tulane University, New Orleans.
- Goes PB, Lin M, Ching-man AY (2014) “Popularity effect” in user-generated content: Evidence from online product reviews. *Inform. Systems Res.* 25(2):222–238.
- Good D, Michel EJ (2013) Individual ambidexterity: Exploring and exploiting in dynamic contexts. *J. Psych.* 147(5):435–453.
- Goodman-Delahunty J, Granhag PA, Hartwig M, Loftus EF (2010) Insightful or wishful: Lawyers' ability to predict case outcomes. *Psych. Public Policy Law* 16(2):133–157.
- Gopal A, Sivaramakrishnan K, Krishnan MS, Mukhopadhyay T (2003) Contracts in offshore software development: An empirical analysis. *Management Sci.* 49(12):1671–1683.
- Grewal D, Gotlieb J, Marmorstein H (1994) The moderating effects of message framing and source credibility on the price-perceived risk relationship. *J. Consumer Res.* 21(1):145–153.
- Grinold R, Kahn R (2000) *Active Portfolio Management* (McGraw Hill, New York).
- Guadagni PM, Little JDC (1983) A logit model of brand choice calibrated on scanner data. *Marketing Sci.* 2(3):203–238.
- Heckman JJ (1979) Sample selection bias as a specification error. *Econometrica* 47(1):153–161.
- Holmqvist M (2004) Experiential learning processes of exploitation and exploration within and between organizations: An empirical study of product development. *Organ. Sci.* 15(1):70–81.
- Hong Y, Pavlou PA (2017) On buyer selection of service providers in online outsourcing platforms for IT services. *Inform. Systems Res.* 28(3):547–562.
- Hong Y, Wang C, Pavlou PA (2015) Comparing open and sealed bid auctions: Evidence from online labor markets. *Inform. Systems Res.* 27(1):49–69.
- Hong Y, Huang N, Burtch G, Li C (2016) Culture, conformity, and emotional suppression in online reviews. *J. Assoc. Inform. Systems* 17(11). <https://aisel.aisnet.org/jais/vol17/iss11/2/#:~:text=We%20found%20that%20consumers%20from,emotions%20to%20be%20less%20helpful>.
- Horton JJ (2010) Online labor markets. Saberi A, ed. *Internet and Network Economics*, Lecture Notes in Computer Science, vol. 6484 (Springer, Berlin Heidelberg).
- Horton JJ, Chilton LB (2010) The labor economics of paid crowdsourcing. *Proc. 11th ACM Conf. Electronic Commerce* (ACM, New York), 209–218.
- Hull CL (1930) Simple trial and error learning: A study in psychological theory. *Psych. Rev.* 37(3):241–256.
- Ipeirotis PG, Provost F, Wang J (2010) Quality management on Amazon Mechanical Turk. *Proc. ACM SIGKDD Workshop Human Comput.* (ACM, New York), 64–67.
- Jain H, Padmanabhan B, Pavlou PA, Santanam RT (2018) Call for papers—Special issue of *Information Systems Research*—Humans, algorithms, and augmented intelligence: The future of work, organizations, and society. *Inform. Systems Res.* 29(1):250–251.
- Jerath K, Fader PS, Hardie GSB (2011) New perspectives on customer “death” using a generalization of the Pareto/NBD model. *Marketing Sci.* 30(5):866–880.
- Jueptner M, Frith CD, Brooks DJ, Frackowiak R, Passingham R (1997) Anatomy of motor learning. II. Subcortical structures and learning by trial and error. *J. Neurophysiology* 77(3):1325–1337.
- Kahneman D, Lovallo D (1993) Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Sci.* 39(1):17–31.
- Kokkodis M (2018) Dynamic recommendations for sequential hiring decisions in online labor markets. *Internat. Conf. Knowledge Discovery Data Mining* (ACM, New York), 453–461.
- Kokkodis M (2021) Dynamic, multi-dimensional, and skillset-specific reputation systems for online work. *Inform. Systems Res.* 32(3):688–712.

- Kokkodis M, Ipeirotis PG (2014) The utility of skills in online labor markets. *Internat. Conf. Inform. Systems*. <https://aisel.aisnet.org/icis2014/proceedings/DecisionAnalytics/15/>.
- Kokkodis M, Ipeirotis PG (2016) Reputation transferability in online labor markets. *Management Sci.* 62(6):1687–1706.
- Kokkodis M, Papadimitriou P, Ipeirotis PG (2015) Hiring behavior models for online labor markets. *Proc. Eighth ACM Internat. Conf. Web Search Data Mining* (ACM, New York), 223–232.
- Kossinets G, Watts DJ (2009) Origins of homophily in an evolving social network. *Amer. J. Sociol.* 115(2):405–450.
- Kreye ME, Yee MG, Newnes LB, Goodwin P (2012) Approaches to displaying information to assist decisions under uncertainty. *Omega* 40(6):682–692.
- Lauren D (2017) Snagajob appoints former Upwork CEO to board of directors. Accessed November 28, 2020, <https://www.prnews.wire.com/news-releases/snagajob-appoints-former-upwork-ceo-to-board-of-directors-300417689.html>.
- Lee S (2019) Learning-by-moving: Can reconfiguring spatial proximity between organizational members promote individual-level exploration? *Organ. Sci.* 30(3):467–488.
- Lee S, Meyer-Doyle P (2017) How performance incentives shape individual exploration and exploitation: Evidence from microdata. *Organ. Sci.* 28(1):19–38.
- Levinthal DA, March JG (1993) The myopia of learning. *Strategic Management J.* 14(S2):95–112.
- Liang C, Hong Y, Gu B (2017) Moral hazards and effects of IT-enabled monitoring systems in online labor markets. *Internat. Conf. Inform. Systems* (Association for Information Systems, Atlanta), 1–10.
- Lin M, Viswanathan S (2015) Home bias in online investments: An empirical study of an online crowdfunding market. *Management Sci.* 62(5):1393–1414.
- Lin M, Liu Y, Viswanathan S (2016) Effectiveness of reputation in contracting for customized production: Evidence from online labor markets. *Management Sci.* 64(1):345–359.
- Liu Y, Huang X, An A, Yu X (2008) Modeling and predicting the helpfulness of online reviews. Giannotti F, Gunopulos D, Turini F, Zaniolo C, Ramakrishnan, Wu X, eds. *IEEE Internat. Conf. Data Mining* (IEEE, Piscataway, NJ), 443–452.
- LIWC (2018) Linguistic inquiry and word count. Accessed November 28, 2020. <http://liwc.wpengine.com/>.
- Lu Y, Tsaparas P, Ntoulas A, Polanyi L (2010) Exploiting social context for review quality prediction. *Proc. 19th Internat. Conf. World Wide Web* (ACM, New York), 691–700.
- March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.
- Mason W, Watts DJ (2010) Financial incentives and the performance of crowds. *ACM SIGKDD Explorations Newsletter* 11(2):100–108.
- McFadden D (1973) Conditional logit analysis of qualitative choice behavior. Zarembka P, ed. *Frontiers in Econometrics* (Academic Press, New York), 105–142.
- Menkhoff L, Schmidt U, Brozynski T (2006) The impact of experience on risk taking, overconfidence, and herding of fund managers: Complementary survey evidence. *Eur. Econom. Rev.* 50(7):1753–1766.
- Mom TJM, Van Den Bosch FAJ, Volberda HW (2007) Investigating managers' exploration and exploitation activities: The influence of top-down, bottom-up, and horizontal knowledge inflows. *J. Management Stud.* 44(6):910–931.
- Mom TJM, Van Den Bosch FAJ, Volberda HW (2009) Understanding variation in managers' ambidexterity: Investigating direct and interaction effects of formal structural and personal coordination mechanisms. *Organ. Sci.* 20(4):812–828.
- Moreno A, Terwiesch C (2014) Doing business with strangers: Reputation in online service marketplaces. *Inform. Systems Res.* 25(4):865–886.
- Nelson P (1970) Information and consumer behavior. *J. Political Econom.* 78(2):311–329.
- Palich LE, Bagby DR (1995) Using cognitive theory to explain entrepreneurial risk-taking: Challenging conventional wisdom. *J. Bus. Venturing* 10(6):425–438.
- Pallais A (2014) Inefficient hiring in entry-level labor markets. *Amer. Econom. Rev.* 104(11):3565–3599.
- Peterson DK, Pitz GF (1988) Confidence, uncertainty, and the use of information. *J. Experiment. Psych. Learn. Memory Cognition* 14(1):85–92.
- Plassmann H, O'Doherty J, Shiv B, Rangel A (2008) Marketing actions can modulate neural representations of experienced pleasantness. *Proc. Natl. Acad. Sci. USA* 105(3):1050–1054.
- Rahman H (2018) Don't worship the stars: Ratings inflation in online labor markets. Pries-Heje J, Ram S, Rosemann M, eds. *Proc. Internat. Conf. Inform. Systems* (Association for Information Systems, Atlanta).
- Rao AR, Monroe KB (1989) The effect of price, brand name, and store name on buyers' perceptions of product quality: An integrative review. *J. Marketing Res.* 26(3):351–357.
- Robinson AT, Marino LD (2015) Overconfidence and risk perceptions: Do they really matter for venture creation decisions? *Internat. Entrepreneurship Management J.* 11(1):149–168.
- Rogan M, Mors ML (2014) A network perspective on individual-level ambidexterity in organizations. *Organ. Sci.* 25(6):1860–1877.
- Rosing K, Zacher H (2017) Individual ambidexterity: The duality of exploration and exploitation and its relationship with innovative performance. *Eur. J. Work Organ. Psych.* 26(5):694–709.
- Schilling MA, Vidal P, Ployhart RE, Marangoni A (2003) Learning by doing something else: Variation, relatedness, and the learning curve. *Management Sci.* 49(1):39–56.
- Sheng VS, Provost F, Ipeirotis PG (2008) Get another label? Improving data quality and data mining using multiple, noisy labelers. *Internat. Conf. Knowledge Discovery Data Mining* (ACM, New York), 614–622.
- Shiv B, Carmon Z, Ariely D (2005) Placebo effects of marketing actions: Consumers may get what they pay for. *J. Marketing Res.* 42(4):383–393.
- Skinner BF (1990) *The Behavior of Organisms: An Experimental Analysis* (BF Skinner Foundation, Cambridge, MA).
- Smith-Keiling BL, Hyun HIF (2019) Applying a computer-assisted tool for semantic analysis of writing: Uses for STEM and ELL. *J. Microbiology Biol. Ed.* 20(1).
- Snir EM, Hitt LM (2003) Costly bidding in online markets for IT services. *Management Sci.* 49(11):1504–1520.
- Snyder CR, Lopez SJ (2009) *Oxford Handbook of Positive Psychology*, Oxford Library of Psychology (Oxford University Press, Oxford, UK).
- Sridhar S, Srinivasan R (2012) Social influence effects in online product ratings. *J. Marketing* 76(5):70–88.
- Sundararajan A (2016) *The Sharing Economy: The End of Employment and the Rise of Crowd-Based Capitalism* (MIT Press, Cambridge, MA).
- Svveaney LC (1999) The role of perceived risk in the quality-value relationship: A study in a retail environment. *J. Retailing* 75(1):77–105.
- Tripp MT, Grégoire Y (2011) When unhappy customers strike back on the internet. *MIT Sloan Management Rev.* 52(3):37–44.
- Wooldridge MJ (2010) *Econometric Analysis of Cross Section and Panel Data* (MIT Press, Cambridge, MA).
- Yelle LE (1979) The learning curve: Historical review and comprehensive survey. *Decision Sci.* 10(2):302–328.
- Yin D, Bond S, Zhang H (2014) Anxious or angry? Effects of discrete emotions on the perceived helpfulness of online reviews. *MIS Quart.* 38(2):539–560.
- Yoganarasimhan H (2013) The value of reputation in an online free-lance marketplace. *Marketing Sci.* 32(6):860–891.
- Zheng A, Hong Y, Pavlou PA (2016) Matching in two-sided platforms for IT services: Evidence from online labor markets. Working paper, Temple University, Philadelphia.