

Estimation of a Non-Parametric Principal-Agent Model with Hidden Actions

Nur Kaynar

UCLA Anderson School of Management, sema.nur.kaynar.keles.phd@anderson.ucla.edu.

Auyon Siddiq

UCLA Anderson School of Management, auyon.siddiq@anderson.ucla.edu.

The design of performance-based incentives can be naturally posed as a moral hazard principal-agent problem. In this setting, a key input to the principal’s optimal contracting problem is the agent’s production function – the dependence of agent output on effort. While agent production is classically assumed to be known to the principal, this is unlikely to be the case in practice. Motivated by the design of performance-based incentives, we present a method for estimating a principal-agent model from data on incentive contracts and associated outcomes, with a focus on estimating agent production. The proposed estimator is statistically consistent and can be expressed as a mathematical program. To circumvent computational challenges with solving the estimation problem exactly, we approximate it as an integer program, which we solve through a column generation algorithm that uses hypothesis tests to select variables. We show that our approximation scheme and solution technique both preserve the estimator’s consistency and combine to dramatically reduce the computational time required to obtain sound estimates. To demonstrate our method, we conducted an experiment on a crowdwork platform (Amazon Mechanical Turk) by randomly assigning incentive contracts with varying pay rates among a pool of workers completing the same task. We present numerical results illustrating how our estimator combined with experimentation can shed light on the efficacy of performance-based incentives.

1. Introduction

The extent to which financial incentives increase worker performance is of interest in many employment settings. This question has taken on renewed relevance due to the emergence of online labor platforms, which are used for on-demand jobs like ride-hailing (e.g., Uber, Lyft), delivery (Postmates), freelance work (Upwork), and short, discrete tasks (Amazon Mechanical Turk). Although these platforms support different types of work, they also have common features: workers are hired and compensated on a per-task basis, work is done remotely with limited supervision, and workers may be offered performance-based incentives.¹

¹ For example, Lyft offers drivers bonuses for fulfilling a target number of rides within a predefined time frame (Lyft 2021), and Postmates offers a similar incentive (Postmates 2021). Similarly, freelance platforms Upwork and Amazon Mechanical Turk allow clients to provide workers with bonuses at their own discretion.

The design of performance-based incentives can be naturally posed as a moral-hazard principal-agent problem, in which an agent’s (worker’s) effort is hidden from the principal (employer), and the agent’s output depends stochastically on their effort (Holmstrom 1979, Grossman and Hart 1983, Sappington 1991). In this setting, the relationship between worker output and effort corresponds to a set of parameters that define agent production. If these parameters are known, then the principal’s problem of optimally designing incentives is well-defined and potentially convex (Grossman and Hart 1983).

In practice, however, the relationship between worker effort and output is unlikely to be known *a priori*. Given data on incentives and associated output, this dependence can be inferred by specifying an appropriate agent model and estimating the parameters that govern agent production. Despite the importance of principal-agent models to the analysis of incentive contracts, estimation problems of this nature are scarce in the literature, even for simple agent models. Estimating an agent model from observational or experimental data can be a useful step toward the design of incentive contracts in practice, and can also play a role in estimating agent welfare under a given contract.

Our main contribution is to present an estimator for a principal-agent model with hidden actions, along with an algorithm for solving the estimation problem. Our focus is on estimating model parameters that encode agent production, namely, the conditional distribution over output for each effort level. To reflect a moral-hazard setting, we assume no data is available on agent effort, which makes the estimation problem computationally non-trivial. We make two methodological contributions in particular: (1) we provide an estimator that is statistically *consistent* under appropriate conditions, meaning it uncovers the true model parameters as the sample size goes to infinity, and (2) we develop an accompanying solution technique that is computationally efficient and preserves consistency.

The agent model we consider is non-parametric, in that we do not assume functional forms for the dependence of agent output on effort, and we assume both output and effort levels are discrete. This specification has two important consequences. First, it admits a simple and *tractable* formulation of a general optimal contracting problem, which allows us to readily solve for an optimal contract under the estimated agent model. Second, estimating agent models is well-known to be challenging due to a need to embed the agent’s problem – itself an optimization problem – within the estimator (Bajari et al. 2007). Our modeling approach allows us to express the estimator as an integer program, which admits a structure that supports obtaining estimates quickly using a novel solution technique. In addition to these computational advantages, our non-parametric model naturally handles threshold-based incentives, which commonly arise in practice, and is flexible enough to have strong predictive performance on a variety of datasets without overfitting.

In an empirical study, we show how our estimator can be combined with experimentation to characterize worker output over a class of incentive contracts, which in turn allows us to solve for an optimal contract from the given class. In a randomized experiment, we recruited a pool of 500 workers from a crowdwork platform (Amazon Mechanical Turk), each of whom was asked to complete an identical proofreading task, with output measured by the number of typos identified. We created exogenous variation in payments by randomly generating the parameters of an incentive contract for each worker. We then applied our estimator to the experimental data to investigate the effect of performance-based incentives on worker output. Our results complement existing findings that incentives do increase output in crowdwork, although we observe diminishing returns to output beginning at relatively low payments.

Our model has limitations. The agent model does not include common features of principal-agent problems; in particular, we do not address risk aversion or unobserved agent heterogeneity in this paper. This abstraction arises from our focus on obtaining consistent estimates (potentially for a large number of parameters) while maintaining computational tractability. Generalizing our estimation procedure to accommodate a richer class of agent models may expand its applicability in practice. Further, our non-parametric approach may be unsuitable for settings with limited data, because it may require estimating many parameters if the action or outcome space is large.

The remainder of the paper is organized as follows. §2 defines the agent model, presents the associated estimator, and establishes consistency. §3 presents an exact formulation of the estimator as an integer program and discusses the computational challenges of the exact representation. §4 develops an approximate estimator and an accompanying solution technique, which dramatically improve tractability while preserving consistency of the exact estimator. §5 describes the randomized experiment and demonstrates the application of our estimator to experimental data. §6 concludes. All proofs are contained in the electronic companion.

1.1. Related literature

Existing work on estimating principal-agent models is relatively limited. Several papers have focused on employee compensation. Ferrall and Shearer (1999) use payroll records of copper mine workers to estimate the cost of employee risk aversion. Paarsch and Shearer (2000) use a tree-planting firm’s records to estimate the impact of providing piece-rate compensation over fixed wages, and Shearer (2004) addresses the same question through a field experiment. Duflo et al. (2012) estimate an agent model to assess the impact of financial incentives for schoolteachers and use the model to estimate cost reductions associated with a counterfactual payment scheme. Misra et al. (2005) and Misra and Nair (2011) both estimate agent models based on salesforce compensation and empirically validate the models on out-of-sample data. Gayle and Miller (2015) focus

on identifying a general principal-agent model motivated by managerial compensation. Georgiadis and Powell (2021) provide conditions under which a single A/B test can estimate the impact of marginal changes to an incentive contract, using the classical principal-agent model from Holmstrom (1979). Applications beyond employee compensation include agriculture (de Zegher et al. 2017) and healthcare (Vera-Hernandez 2003, Lee and Zenios 2012, Aswani et al. 2019).

Previous work on estimating principal-agent models have used a variety of methods, including least squares (Lee and Zenios 2012), simulated method of moments (Paarsch and Shearer 2000, Misra et al. 2005, Misra and Nair 2011, Duflo et al. 2012), simulation-based maximum likelihood estimation (Ferrall and Shearer 1999, Vera-Hernandez 2003, Aswani et al. 2019), and numerical minimization of a sum-of-squares criterion (Gayle and Miller 2015). Our approach differs in that we formulate the estimation problem as an integer program, which is made possible by our specification of the agent model, in particular by assuming agent actions and outputs are discrete.

We solve the estimation problem using a column generation algorithm that exploits statistical properties of the formulation. Column generation methods have been successfully applied to solve large-scale linear and integer programs in which an extremely large number of variables is the main obstacle to obtaining optimal solutions (Vanderbeck and Wolsey 1996, Barnhart et al. 1998, Lubbecke and Desrosiers 2005). These methods typically involve solving a tractable *master* problem that restricts attention to a subset of decision variables, and selectively introducing variables into the formulation until a certificate of optimality or alternative termination criterion is met. In contrast to existing column generation methods that select columns using dual information, our algorithm uses a series of non-parametric hypothesis tests to identify variables to introduce into the master problem. This approach is viable in our setting because the decision variables are mapped to empirical probability distributions constructed from the data, giving them a clear statistical interpretation. By comparison, existing column generation methods have typically been applied to deterministic settings where the model parameters may not have any statistical meaning (see Lubbecke and Desrosiers (2005) for a review).

The estimation problem we consider is also closely related to a recent line of research on *inverse* optimization, in which optimization model parameters are inferred from (potentially noisy) solution data. Existing approaches to inverse optimization have focused on estimating parameters of linear programs (Chan et al. 2018) or general convex optimization problems (Keshavarz et al. 2011, Bertsimas et al. 2015, Aswani et al. 2018). Similar to our paper, the literature on inverse optimization is often motivated by an interest in estimating a model of agent decision-making from data (Aswani et al. 2018, Esfahani et al. 2018). Our paper differs in that instead of assuming the agent solves a convex optimization problem, we assume they select a utility-maximizing action from a discrete set, which calls for a different solution approach.

1.2. Notation

For convenience, we describe notational conventions here. Sets are denoted by upper case letters, scalars by lower case letters, and vectors and matrices by lower case, boldfaced letters. For a $m \times d$ matrix \mathbf{x} , let \mathbf{x}_a be the a^{th} row, and let x_{aj} be the entry in the a^{th} row and j^{th} column. For vectors \mathbf{x} and \mathbf{y} , let $\|\mathbf{x}\|_1 = \sum_{a=1}^m \sum_{j=1}^d |x_{aj}|$ denote the ℓ_1 -norm, and let $\mathbf{x} \circ \mathbf{y} = \sum_{a=1}^m \sum_{j=1}^d x_{aj} y_{aj}$ be the elementwise product. For a matrix of random variables \mathbf{x}_n , we use both $\mathbf{x}_n \rightarrow \mathbf{x}^0$ and $\text{plim}_{n \rightarrow \infty} \mathbf{x}_n = \mathbf{x}^0$ to mean \mathbf{x}_n converges elementwise in probability to \mathbf{x}^0 as $n \rightarrow \infty$, unless otherwise specified. Define the indicator variable $\mathbb{I}\{\cdot\} = 1$ if the statement $\{\cdot\}$ is true, and 0 otherwise. For simplicity, we use $\mathbb{E}(\cdot)$ for all expectations and $\text{Pr}(\cdot)$ for all probabilities throughout the paper.

2. Estimator

In this section, we define the principal-agent model (§2.1), formulate the estimator (§2.2), and prove its statistical consistency (§2.3).

2.1. Principal-agent model and contract data

Our principal-agent model is a discrete analogue to the canonical model introduced by Grossman and Hart (1983). We choose this model for both its simplicity and generality. The interaction between the principal and agent proceeds as follows. The principal selects a *contract* to offer the agent, which is a mapping of payments to *outcomes* (i.e., agent output). Outcomes depend stochastically on a costly *action* (i.e., effort) taken by the agent. Outcomes are observed by both parties, while actions are observed only by the agent.

Let A and J index the set of possible actions and outcomes, respectively, where $|A| = m$ and $|J| = d$. Let ξ be a discrete random variable denoting the outcome, where $\xi \in J$. We denote a contract by $\mathbf{r} \in \mathbb{R}_+^d$, where r_j is the payment to the agent if outcome j is realized. Let $\mathbf{c} \in \mathbb{R}_+^m$ denote action costs, where c_a is the cost to the agent of taking action a . The dependence of outcomes on actions is governed by a parameter matrix $\boldsymbol{\pi} \in \mathbb{R}_+^{m \times d}$, where π_{aj} denotes the probability that action a leads to outcome j . We use $\boldsymbol{\pi}_a \in \mathbb{R}^d$ to denote the probability mass function over outcomes associated with action a .

Given a contract \mathbf{r} , the agent selects an action to maximize their expected utility by solving

$$\max_{a \in A} \left\{ \sum_{j \in J} \pi_{aj} r_j - c_a \right\}. \quad (1)$$

We assume that there exists at least one action that yields non-negative expected utility for the agent. If for each $a \in A$ the distribution $\boldsymbol{\pi}_a$ is known, the principal's problem of selecting a utility-maximizing contract can be formulated as a convex optimization problem (Grossman and Hart 1983). We take an inverted perspective in this paper, by instead supposing that the distributions

π_a , $a \in A$ are unknown, but may be estimated given appropriate data. In particular, suppose we have data from n identical agents,²

$$(\mathbf{r}^i, \xi^i), i \in I, \quad (2)$$

where I indexes pairs of incentive contracts and outcomes, and $|I| = n$. Let $R \subseteq \mathbb{R}^d$ be the set of all possible values of \mathbf{r}^i . Further, we assume the contract set R is bounded, in that there exists a constant \bar{r} such that $\bar{r} = \sup_{\mathbf{r} \in R} \|\mathbf{r}\|_0 < \infty$. The assumption that R is bounded ensures that the contracts \mathbf{r}^i remain bounded as $n \rightarrow \infty$.

Next, suppose we have no observations of past agent actions, and only know the agent's action set A and associated costs, \mathbf{c} . A natural question in this setting is to predict the distribution of the outcome ξ^{n+1} under a new contract \mathbf{r}^{n+1} . Note that if π is known, then this prediction task reduces to solving the agent's problem (1) under \mathbf{r}^{n+1} , identifying the optimal action a , and taking π_a to be the distribution of ξ^{n+1} . Therefore, the matrix π is the key model primitive for predicting the outcome associated with \mathbf{r}^{n+1} . Our goal is to estimate the parameter π from data that takes the form given in (2).

The assumption that agent costs are unknown is relatively mild in our setting, given that agent actions are also hidden. From a model-fitting perspective, it suffices to select \mathbf{c} to cover a range of *possible* costs to the agent. In our numerical study in §5, we take a machine learning perspective by treating the number of agent actions m and the set of costs \mathbf{c} as hyperparameters that are tuned prior to fitting the model.

2.2. Estimator formulation

Next, we formalize the estimator for π .³ Let

$$A(\mathbf{r}, \pi) = \operatorname{argmax}_{a \in A} \left\{ \sum_{j \in J} \pi_{aj} r_j - c_a \right\} \quad (3)$$

denote the set of optimal actions under the contract \mathbf{r} and the model π . Let $\mathbf{y} \in \{0, 1\}^{n \times d}$ be a binary matrix that encodes historical outcomes, where $y_j^i = 1$ if $\xi^i = j$ and $y_j^i = 0$ if $\xi^i \neq j$. For each $i \in I$, let x^i be a decision variable representing the agent action under contract \mathbf{r}^i , and let $\omega \in \mathbb{R}_+^{m \times d}$ be a set of auxiliary variables, which will be used to model empirical probabilities. For fixed π , the loss function $L_n(\pi)$ is then given by

$$L_n(\pi) = \operatorname{minimize}_{\mathbf{x}, \omega} \sum_{a \in A} \sum_{j \in J} |\pi_{aj} - \omega_{aj}| \quad (4a)$$

² We extend our model to accommodate heterogeneous agents in §EC.3 of the electronic companion.

³ Throughout the paper, we shall use *estimator* to refer to an optimization problem or algorithm, and *estimate* to refer to its solutions.

$$\text{subject to } x^i \in A(\mathbf{r}^i, \boldsymbol{\pi}), \quad i \in I, \quad (4b)$$

$$\omega_{aj} = \frac{1}{|\{i|x^i = a\}|} \sum_{i \in \{i|x^i = a\}} y_j^i, \quad a \in A, j \in J. \quad (4c)$$

In the formulation above, (4b) restricts each x^i to be an optimal action under \mathbf{r}^i and $\boldsymbol{\pi}$, and (4c) defines ω_{aj} to be the empirical probability that action a leads to outcome j . Note that the empirical probability ω_{aj} depends on the cardinality of the set $\{i|x^i = a\}$, which is the implied number of data points for which the action a is optimal for the agent under $\boldsymbol{\pi}$. The objective (4a) then simply measures the error between the model probabilities $\boldsymbol{\pi}$ and the implied empirical probabilities $\boldsymbol{\omega}$.

Next, let Π be a compact set representing the parameter set for $\boldsymbol{\pi}$. The estimate is then attained at a minimizer of the loss function over Π :

$$(PA) \quad \hat{\boldsymbol{\pi}}_n \in \underset{\boldsymbol{\pi} \in \Pi}{\operatorname{argmin}} L_n(\boldsymbol{\pi}).$$

It will be convenient to interpret the parameter set Π as the Cartesian product of m probability simplices – one for each action $a \in A$.

2.3. Statistical consistency

Let us now suppose there exists a “true” model parameter $\boldsymbol{\pi}^0$ that is responsible for generating the data (\mathbf{r}^i, ξ^i) , $i \in I$. We say an estimator is *statistically consistent* if it produces a sequence of estimates $\hat{\boldsymbol{\pi}}_n$ such that $\hat{\boldsymbol{\pi}}_n \rightarrow \boldsymbol{\pi}^0$ as $n \rightarrow \infty$. This raises a natural question: Under what conditions, if any, is PA a consistent estimator? In general, whether an estimator is consistent depends on the specification of the loss function. Our main result in this section, Theorem 1, shows that minimizing the loss function $L_n(\boldsymbol{\pi})$ defined in (4) produces an estimate that is indeed consistent.

Before addressing the consistency of PA, we first formalize the statistical model that generates the data. First, we define an important set that is used throughout our analysis:

$$R_a(\boldsymbol{\pi}) = \left\{ \mathbf{r} \in R \mid a \in \underset{a \in A}{\operatorname{argmax}} \sum_{j \in J} \pi_{aj} r_j - c_a \right\}, \quad (5)$$

Here, $R_a(\boldsymbol{\pi})$ represents the subset of the contract set R where action $a \in A$ is optimal for the agent, given the model $\boldsymbol{\pi}$. Next, we impose two assumptions. The first assumption formalizes the data generation process.

ASSUMPTION 1 (Data). *The data (\mathbf{r}^i, ξ^i) , $i \in I$, are independent samples of random variables (\mathbf{r}, ξ) , where (i) (\mathbf{r}, ξ) are jointly distributed with support $R \times J$, (ii) \mathbf{r} has continuous marginal density function $f(\mathbf{r})$, (iii) $\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi})) > 0$ for all $a \in A$ and $\boldsymbol{\pi} \in \Pi$, and (iv) ξ has conditional mass function $\pi_{aj}^0 = \Pr(\xi = j | \mathbf{r} \in R_a(\boldsymbol{\pi}^0))$, where $\boldsymbol{\pi}^0 \in \Pi$.*

Assumption 1(*iv*) states that there exists a “true” parameter – denoted $\boldsymbol{\pi}^0$ – that is responsible for generating the outcomes ξ^i , based on the agent model (1). The statements in (*ii*) and (*iii*) are regularity conditions that we use to prove convergence of $\hat{\boldsymbol{\pi}}_n$ to $\boldsymbol{\pi}^0$.⁴ Our assumption that the data are independent and identically distributed (i.i.d.) is commonly used in the statistical learning literature to obtain similar consistency results.⁵

Next, we consider an additional condition that is important for our main result in Theorem 1.

ASSUMPTION 2 (Identifiability). *For every $\boldsymbol{\pi} \in \Pi$ such that $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$, there exists an (a, j) such that*

$$\pi_{aj} \neq \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^0) | \mathbf{r} \in R_a(\boldsymbol{\pi})).$$

Assumption 2 is an *identifiability* condition, which ensures that the unknown parameter $\boldsymbol{\pi}^0$ can be learned from the data. This assumption implies a one-to-one mapping between the parameter set Π and the joint distribution of the random variables (\mathbf{r}, ξ) . In other words, Assumption 2 guarantees that the distribution of (\mathbf{r}, ξ) is unique for each $\boldsymbol{\pi} \in \Pi$. In the absence of model identifiability, there may exist multiple parameters values in Π that generate the same distribution in the observed data; in this case, it is impossible for any estimation procedure to pinpoint the true $\boldsymbol{\pi}^0$. Identifiability conditions like Assumption 2 are commonly imposed to prove consistency of an estimator (Van der Vaart 2000).

We can now present the main result of §2, which shows that the estimator PA uncovers the true model parameter $\boldsymbol{\pi}^0$ under Assumptions 1 and 2.

THEOREM 1. *Let Assumption 1 hold. Then $\hat{\boldsymbol{\pi}}_n \rightarrow \boldsymbol{\pi}^0$ for any $\boldsymbol{\pi}^0 \in \Pi$ if and only if Assumption 2 holds.*

Theorem 1 states that the estimator PA is statistically *consistent*, which is defined as the convergence of estimates to the true model parameters (Van der Vaart 2000, Casella and Berger 2002, Bickel and Doksum 2015). Despite being an asymptotic property, consistency is valuable in practice, because it guarantees that parameter estimates will generally improve with additional data.

⁴ The assumption that the contract data \mathbf{r}^i , $i \in I$ is generated by a continuous density function $f(\mathbf{r})$ is important for our technical results. Intuitively, because the \mathbf{r}^i are input data, assuming this continuity provides the estimator with more information, which makes precise inference of $\boldsymbol{\pi}^0$ possible under the identifiability condition in Assumption 2. If the contract data is instead generated by a discrete distribution supported on a subset of R , then a stronger identifiability than Assumption 2 is needed to compensate for the loss of information. We consider such a case in §EC.4 of the electronic companion.

⁵ Because we assume the data is generated by n independent agents making decisions simultaneously, which is plausible in online labor platforms, the i.i.d. assumption is not particularly restrictive for our setting. Moreover, this assumption is not strictly necessary to achieve consistency, depending on the problem setup. In §EC.4 of the electronic companion, we consider a variation of the model where $\boldsymbol{\pi}^0$ can be estimated by dynamically selecting the contracts to offer the agent. This breaks the independence assumption on the contracts \mathbf{r}^i , but allows for consistent estimation of $\boldsymbol{\pi}^0$ under a different set of assumptions.

Conversely, an inconsistent estimator may produce inaccurate estimates of the unknown parameters, even if data is abundant. In pathological cases, the accuracy of an inconsistent estimator may even decrease with additional data. Therefore, a proof of consistency provides some assurance that parameter estimates will be “reasonable” under moderate sample sizes, and that the accuracy of the estimates will continue to improve with additional data.

Having established that the estimate $\hat{\pi}_n$ behaves desirably, we now shift our attention to solving the estimator PA. Note that in a setting where agent actions are observable, a consistent estimate of π^0 can be obtained by simply counting the relative frequency of outcomes associated with each action. In contrast, when agent actions are hidden, the estimation problem is non-trivial. At a high level, our approach for solving PA will be to leverage integer programming within a broader solution algorithm. The key challenge we face in solving PA is to develop a solution method that satisfies two criteria: (1) is computationally efficient, and (2) preserves the statistical consistency of PA. We note here that an alternative solution approach might be to formulate and solve a convex approximation to PA, although doing so may result in an inconsistent estimator. We will therefore focus on obtaining solutions to PA directly.

3. Exact Integer Programming Formulation

In this section, we present an approach for solving PA exactly using integer programming. We will assume throughout that the parameter set Π is of the form

$$\Pi = \left\{ \pi \in Q_\pi \mid \pi \geq 0, \sum_{j \in J} \pi_{aj} = 1 \text{ for } a \in A \right\}, \quad (6)$$

where Q_π is a polyhedron defined by a set of linear inequalities in π . Assuming that $\pi \in Q_\pi$ permits the formulation of the estimator as a mixed-integer linear program, while also allowing Π to capture various shape constraints on the parameter π . For example, if

$$Q_\pi = \left\{ \pi \mid \sum_{k=j}^d \pi_{ak} \leq \sum_{k=j}^d \pi_{(a+1)k}, a \in \{1, 2, \dots, m-1\}, j \in J \right\}, \quad (7)$$

then for any $a \in \{1, 2, \dots, m-1\}$, Π forces the distribution π_{a+1} to stochastically dominate π_a in the first order, meaning costlier actions taken by the agent are more likely to generate high output. Alternatively, if $Q_\pi = \mathbb{R}^{m \times d}$, then Π permits each π_a to be any valid probability mass function over the outcomes J . We will assume throughout that Π satisfies (6) unless otherwise stated.

Although PA is based on an intuitive loss function, a naive formulation of PA as a mathematical program yields non-linear terms in the objective, due to how the variable ω enters the loss expression (4a). However, the estimation problem is amenable to mathematical programming approaches under a slight modification. Consider the proxy loss function

$$\begin{aligned}
Z_n(\boldsymbol{\pi}) = \underset{\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\omega}}{\text{minimize}} \quad & \frac{1}{n} \sum_{a \in A} \sum_{j \in J} \eta_{aj} |\pi_{aj} - \omega_{aj}| \\
& \eta_{aj} = |\{i | x^i = a\}|, \quad a \in A, j \in J, \\
& (4b) - (4c).
\end{aligned} \tag{8}$$

Here, η_{aj} is the number of observations for which action a is implied to be optimal for the agent under $\boldsymbol{\pi}$. The loss function $Z_n(\boldsymbol{\pi})$ can be interpreted as a scaled version of $L_n(\boldsymbol{\pi})$, where the (a, j) component of $L_n(\boldsymbol{\pi})$ is scaled by η_{aj}/n . The proxy estimator is then given by

$$\boldsymbol{\pi}_n^* = \underset{\boldsymbol{\pi} \in \Pi}{\text{argmin}} Z_n(\boldsymbol{\pi}). \tag{9}$$

Next, we show that (9) can be formulated exactly as a mixed-integer linear program. With a slight abuse of notation, let $\mathbf{x} \in \{0, 1\}^{m \times n}$ be binary variables, where $x_a^i = 1$ if $a \in A(\mathbf{r}^i, \boldsymbol{\pi})$, and $x_a^i = 0$ if $a \notin A(\mathbf{r}^i, \boldsymbol{\pi})$. Introducing the auxiliary variables z_{aj} to linearize the absolute values in the objective of (8) (Bertsimas and Tsitsiklis 1997) yields the following formulation:

$$\begin{aligned}
& \underset{\boldsymbol{\pi}, \mathbf{x}, \mathbf{z}}{\text{minimize}} \quad \sum_{a \in A} \sum_{j \in J} z_{aj} & (10a) \\
\text{subject to} \quad & z_{aj} \geq \frac{1}{n} \sum_{i \in I} (y_j^i - \pi_{aj}) x_a^i \quad a \in A, j \in J, & (10b) \\
& z_{aj} \geq \frac{1}{n} \sum_{i \in I} (\pi_{aj} - y_j^i) x_a^i \quad a \in A, j \in J, & (10c) \\
& \left(\sum_{j \in J} \pi_{aj} r_j^i - c_a \right) x_a^i \geq \left(\sum_{j \in J} \pi_{bj} r_j^i - c_b \right) x_a^i, \quad i \in I, a \in A, b \in A, & (10d) \\
(\text{PA-C}) \quad & \sum_{j \in J} \pi_{aj} = 1, \quad a \in A, & (10e) \\
& \sum_{a \in A} x_a^i = 1, \quad i \in I, & (10f) \\
& x_a^i \in \{0, 1\}, \quad a \in A, & (10g) \\
& \pi_{aj} \geq 0, \quad a \in A, j \in J, & (10h) \\
& \boldsymbol{\pi} \in Q_\pi. & (10i)
\end{aligned}$$

The objective (10a) and constraints (10b)–(10c) represent the error function $\frac{1}{n} \|\boldsymbol{\eta} \circ (\boldsymbol{\pi} - \boldsymbol{\omega})\|_1$ given in (8). Constraint (10d) ensures that $x_a^i = 1$ only if $a \in A(\mathbf{r}^i, \boldsymbol{\pi})$, that is, only if action a is optimal under contract \mathbf{r}^i and the parameter $\boldsymbol{\pi}$. Constraint (10e) ensures that the probability vector $\boldsymbol{\pi}_a$ sums to 1 for each $a \in A$, and constraint (10f) forces exactly one action to be selected as optimal for each contract $i \in I$. Next, we establish an equivalence between the proxy estimator PA-C and the original estimator PA.

PROPOSITION 1. *The estimate $\boldsymbol{\pi}_n^*$ attained at a solution to PA-C is (i) a minimizer of the proxy loss function $Z_n(\boldsymbol{\pi})$, (ii) an asymptotic minimizer of the loss function $L_n(\boldsymbol{\pi})$, $|L_n(\boldsymbol{\pi}_n^*) - L_n(\hat{\boldsymbol{\pi}}_n)| \rightarrow 0$, and (iii) consistent, $\boldsymbol{\pi}_n^* \rightarrow \boldsymbol{\pi}^0$.*

In Proposition 1, (i) establishes that the mathematical program PA-C is equivalent to the proxy estimator (9), (ii) establishes that solving PA-C asymptotically produces an optimal solution to PA, and (iii) confirms that PA-C is also a consistent estimator for $\boldsymbol{\pi}^0$. Based on the equivalence in Proposition 1, we will refer to PA-C as the *exact* estimator in the remainder of the paper.

The intuition behind Proposition 1 is as follows. Note that $Z_n(\boldsymbol{\pi})$ can be interpreted as a re-weighted version of $L_n(\boldsymbol{\pi})$, where for each (a, j) , the term $|\pi_{aj} - \omega_{aj}|$ is multiplied by the weight η_{aj}/n . As $n \rightarrow \infty$, the minimal possible loss for both estimators occurs when $\pi_{aj} = \omega_{aj}$ for all (a, j) . Therefore, minimizing $Z_n(\boldsymbol{\pi})$ also minimizes $L_n(\boldsymbol{\pi})$, in the limit.

Next, note that (10a)–(10d) contains bilinear terms due to the product of the decision variables \mathbf{x} and $\boldsymbol{\pi}$. Because \mathbf{x} and $\boldsymbol{\pi}$ are binary and continuous variables, respectively, these product terms can be linearized exactly using well-known reformulation techniques (Glover 1975, Adams et al. 2004), leading to a mixed-integer linear program. However, a drawback of this approach is that linearizing products of variables is known to yield weak linear programming relaxations (Adams et al. 2004, Luedtke et al. 2012), which can make solving PA-C using off-the-shelf optimization solvers challenging, even for moderately-sized data sets. In the next section, we propose an approximation to PA-C that bypasses the linearization step while remaining statistically well-behaved.

4. Restricted Estimator and Statistical Column Generation

We begin this section by proposing an approximation of PA-C – which we call PA-D – based on replacing the parameter set Π with a discrete subset $\tilde{\Pi}$ (§4.1). We then present a data-driven procedure for constructing the parameter set $\tilde{\Pi}$, and investigate the behavior of the resulting estimates (§4.2). Then, to solve PA-D, we present a column generation algorithm based on hypothesis testing, and show that the algorithm preserves statistical consistency (§4.3). We conclude the section by comparing the numerical performance of the statistical column generation algorithm with off-the-shelf optimization solvers (§4.4).

4.1. Restricted estimator

Our approach to approximately solving PA-C will be to minimize the proxy loss $Z_n(\boldsymbol{\pi})$ over a restricted parameter set $\tilde{\Pi} \subseteq \Pi$ instead of Π . The advantage of this “restricted estimator” is that the agent optimality conditions (10d) can be enforced without introducing bilinear terms into the formulation, which allows us to avoid the computational challenges that often accompany linearization techniques.

Next, we define a set that plays a critical role in our estimation procedure: Let $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{|S|}\} \subseteq \mathbb{R}_+^d$ be a set of vectors indexed by S , where $\sum_{j \in J} v_{sj} = 1$ and $\mathbf{v}_s \geq 0$ for all $s \in S$. We refer to each \mathbf{v}_s as a *candidate distribution*. Next, let the restricted parameter set be defined as

$$\tilde{\Pi} = \left\{ \boldsymbol{\pi} \in \Pi \mid \boldsymbol{\pi}_a \in V \text{ for } a \in A \right\}, \quad (11)$$

and let

$$\tilde{\boldsymbol{\pi}}_n = \underset{\boldsymbol{\pi} \in \tilde{\Pi}}{\operatorname{argmin}} Z_n(\boldsymbol{\pi}) \quad (12)$$

be the associated estimate. For each action $a \in A$, the parameter set $\tilde{\Pi}$ restricts the probability distribution $\boldsymbol{\pi}_a$ to lie in the set of candidate distributions V . We assume throughout that $\tilde{\Pi}$ is non-empty.⁶

Similar to the exact estimator (9), the restricted estimator (12) can also be formulated as a mixed-integer linear program. The intuition behind this formulation is to construct the estimate $\boldsymbol{\pi}$ in a row-wise manner by assigning a candidate distribution in V to each row of $\boldsymbol{\pi}$. To that end, let $\mathbf{w} \in \{0, 1\}^{m \times S}$, $\mathbf{x} \in \{0, 1\}^{n \times S}$ and $\boldsymbol{\phi} \in \{0, 1\}^{n \times m \times S}$ be binary variables with the following interpretations: $w_{as} = 1$ if the candidate distribution \mathbf{v}_s is assigned to be the distribution $\boldsymbol{\pi}_a$, $x_s^i = 1$ if the action assigned to candidate distribution \mathbf{v}_s is optimal under contract \mathbf{r}^i , and $\phi_{as}^i = 1$ if the candidate distribution \mathbf{v}_s is assigned to distribution $\boldsymbol{\pi}_a$ and action a is optimal under \mathbf{r}^i and $\boldsymbol{\pi}$. Similar to PA-C, let $\mathbf{z} \in \mathbb{R}_+^{d \times S}$ be auxiliary variables used to linearize the absolute values in the loss function $Z_n(\boldsymbol{\pi})$. Then the restricted estimator (12) is equivalent to the following mixed-integer linear program:

$$\underset{\boldsymbol{\pi}, \mathbf{w}, \mathbf{x}, \mathbf{z}, \boldsymbol{\phi}}{\operatorname{minimize}} \quad \sum_{s \in S} \sum_{j \in J} z_{sj} \quad (13a)$$

$$\text{subject to} \quad z_{sj} \geq \frac{1}{n} \sum_{i \in I} (y_j^i - v_{sj}) x_s^i, \quad s \in S, j \in J, \quad (13b)$$

$$z_{sj} \geq \frac{1}{n} \sum_{i \in I} (v_{sj} - y_j^i) x_s^i, \quad s \in S, j \in J, \quad (13c)$$

$$\sum_{b \in A} \sum_{s \in S} \left(\sum_{j \in J} v_{sj} r_j^i - c_b \right) \phi_{bs}^i \geq \left(\sum_{j \in J} v_{s'j} r_j^i - c_a \right) w_{at}, \quad i \in I, a \in A, s' \in S, \quad (13d)$$

$$\text{(PA-D)} \quad \sum_{s \in S} w_{as} = 1, \quad a \in A, \quad (13e)$$

⁶ The parameter set $\tilde{\Pi}$ may be empty if the requirement that $\boldsymbol{\pi}_a \in V$ for $a \in A$ conflicts with the requirement that $\boldsymbol{\pi} \in Q_\pi$ from (6). In this case, non-emptiness of $\tilde{\Pi}$ can be guaranteed by projecting the candidate distributions contained in V onto the polyhedron Q_π .

$$\sum_{a \in A} \sum_{s \in S} \phi_{as}^i = 1, \quad i \in I, \quad (13f)$$

$$x_s^i = \sum_{a \in A} \phi_{as}^i, \quad i \in I, s \in S, \quad (13g)$$

$$\phi_{as}^i \leq w_{as}, \quad i \in I, a \in A, s \in S, \quad (13h)$$

$$\pi_{aj} = \sum_{s \in S} w_{as} v_{sj}, \quad a \in A, j \in J, \quad (13i)$$

$$x_s^i \in \{0, 1\}, \quad i \in I, s \in S, \quad (13j)$$

$$w_{as} \in \{0, 1\}, \quad a \in A, s \in S, \quad (13k)$$

$$\phi_{as}^i \in \{0, 1\}, \quad i \in I, a \in A, s \in S, \quad (13l)$$

$$\boldsymbol{\pi} \in Q_{\boldsymbol{\pi}}. \quad (13m)$$

The objective (13a) and constraints (13b)–(13c) together represent the loss function $Z_n(\boldsymbol{\pi})$. Constraint (13d) enforces the agent’s optimality conditions by ensuring that $\phi_{as}^i = 1$ only if candidate distribution \mathbf{v}_s is mapped to $\boldsymbol{\pi}_a$, and if action a is optimal for the agent under \mathbf{r}^i and $\boldsymbol{\pi}$. Constraint (13e) forces exactly one candidate distribution in V to be mapped to each distribution $\boldsymbol{\pi}_a$. Constraint (13f) ensures that only one candidate distribution in V and action $a \in A$ is selected for contract \mathbf{r}^i . Constraint (13g) forces $x_s^i = 1$ if candidate distribution \mathbf{v}_s is mapped to $\boldsymbol{\pi}_a$ and if action a is optimal under \mathbf{r}^i and $\boldsymbol{\pi}$. Constraint (13h) ensures $\phi_{as}^i = 1$ only if \mathbf{v}_s is mapped to $\boldsymbol{\pi}_a$. Constraint (13i) defines $\boldsymbol{\pi}_a$ as the candidate distribution from V that is assigned by \mathbf{w} , and constraint (13m) represents additional shape constraints imposed by the polyhedron $Q_{\boldsymbol{\pi}}$. The key distinction between PA-D and PA-C is that the discrete nature of the parameter set allows the key decision variables $(\mathbf{w}, \mathbf{x}, \boldsymbol{\phi})$ to be binary, which allows us to represent the agent’s optimality conditions in a way that circumvents the need for product terms.

Note that $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$ represents the error in the loss function that arises from solving the restricted estimator PA-D instead of the exact estimator PA-C. Next, we present a random clustering procedure for constructing the set of candidate distributions V , and provide a finite-sample characterization of the error $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$ under the proposed procedure.

4.2. Construction of candidate distributions and finite-sample error

Because PA-C is a consistent estimator of $\boldsymbol{\pi}^0$ (by Proposition 1), we might expect PA-D to also produce a reasonable estimate of $\boldsymbol{\pi}^0$ if the loss function error $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$ is sufficiently small. Additionally, note that $Z_n(\tilde{\boldsymbol{\pi}}_n)$ is the minimal loss when the restricted parameter set $\tilde{\Pi}$ is substituted for Π . As a result, the magnitude of the gap $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$ depends on the restricted parameter set $\tilde{\Pi}$, and by extension, the set of candidate distributions V . Here, we present a method for constructing $\tilde{\Pi}$, based on leveraging the observed data (\mathbf{r}^i, ξ^i) , $i \in I$ to guide the construction of V . Our approach to constructing the candidate distributions V is summarized in Algorithm 1.

Algorithm 1: Sample-based construction of candidate distributions

Input: Data $(\mathbf{r}^i, \xi^i), i \in I$, parameter $\rho > 0$.

1. Randomly sample a subset S from I .
2. **for** each $s \in S$:

$$B_s = \{\mathbf{r} \in R \mid \|\mathbf{r}^s - \mathbf{r}\|_2 \leq \rho\},$$

$$I_s = \{i \in I \mid \mathbf{r}^i \in B_s\}.$$

for each $j \in J$:

$$v_{sj} = \frac{1}{n_s} \sum_{i \in I_s} y_j^i.$$

Output: Candidate distributions $V = \{\mathbf{v}_s \text{ for } s \in S\}$.

Algorithm 1 involves selecting subsets of the contract data, computing the empirical mass function over outcomes for each subset, and designating each of these empirical mass functions as a candidate distribution, \mathbf{v}_s . Note that the s^{th} candidate distribution is based on the outcomes of all contracts \mathbf{r}^i that fall within a ball $B_s \subseteq R$; accordingly, we shall refer to the collection of data points indexed by I_s as the s^{th} cluster. The intuition for constructing the candidate distributions V in this manner is simple: Based on the agent model (1), contracts that are within a small distance of each other are likely to induce the same action from the agent. Therefore, the empirical distribution of outcomes for all contracts that lie within the ball B_s can be assumed to approximate one of the rows of the true parameter matrix $\boldsymbol{\pi}^0$ (although which row it approximates remains unknown). Next, we show that the error $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$ is well behaved if V is constructed using Algorithm 1. We first impose the following assumption.

ASSUMPTION 3 (Clustering condition). For each $a \in A$, there exists $s \in S$ such that $B_s \subseteq R_a(\boldsymbol{\pi}^0)$ and $I_s \neq \emptyset$.

Assumption 3 states that for every action a , Algorithm 1 produces a ball B_s that is entirely inside the subset of the contract set R that induces action a from the agent, $R_a(\boldsymbol{\pi}^0)$. Note that if $B_s \subseteq R_a(\boldsymbol{\pi}^0)$ then every contract in cluster s induces action a from the agent. This implies that \mathbf{v}_s is an empirical distribution sampled from $\boldsymbol{\pi}_a^0$. Therefore, Assumption 3 implies that for each row of $\boldsymbol{\pi}^0$, there exists at least one candidate distribution in V that is constructed by sampling from that row. Note that Assumption 3 is more likely to hold when S in Algorithm 1 is large (because we construct many balls B_s) and ρ is small (because each ball is smaller).

Our next result shows that if Assumption 3 and an additional condition on Π holds, we can bound the approximation error $Z_n(\tilde{\boldsymbol{\pi}}_n) - Z_n(\boldsymbol{\pi}_n^*)$.

THEOREM 2. Let Assumption 3 hold, and let V be constructed using Algorithm 1. Further, suppose $\Pi = \{\boldsymbol{\pi} \geq 0 \mid \sum_{j \in J} \pi_{aj} = 1, a \in A\}$. Then there exists $\kappa \in (0, 1)$ such that for any $\varepsilon \in (0, 1)$,

$$\Pr(|Z_n(\boldsymbol{\pi}_n^*) - Z_n(\tilde{\boldsymbol{\pi}}_n)| > \varepsilon) \leq O(n^2 \kappa^n). \quad (14)$$

We offer a few remarks on Theorem 2. First, observe that the bound is not monotonic due to the n^2 term, which implies that the bound can become looser in n for small n . This occurs because our proof approach depends on constructing a feasible solution $\tilde{\pi}$, and bounding the absolute number of observations where the hidden agent action is “misclassified” by $\tilde{\pi}$. Thus, the n^2 term reflects the possibility that the number of misclassified actions may increase with the sample size. Further, note that if κ is close to 1 and n is of moderate size, the bound in Theorem 2 may be vacuous. However, because it is guaranteed that $\kappa \in (0, 1)$, it is straightforward to verify that $n^2\kappa^n \rightarrow 0$, which implies that the error $Z_n(\tilde{\pi}_n) - Z_n(\pi_n^*)$ eventually vanishes in n .

Second, note that the rate depends on the constant κ , with lower values of κ leading to faster convergence. While κ is not particularly interpretable, it can be shown to decrease in ρ and increase in the number of clusters $|S|$. Note from Algorithm 1 that ρ is the radius of the ball B_s , for each cluster $s \in S$. Intuitively, for fixed n , larger values of ρ makes each ball B_s contain a larger number of observations, which leads to faster convergence. Conversely, larger values of $|S|$ will slow convergence, for the following reason: Because the bound depends in part on the cluster that has the fewest observations, large values of $|S|$ will increase the probability that at least one of the clusters has very few data points, which weakens the bound. Therefore, the rate $n^2\kappa^n$ is fastest when ρ is large and $|S|$ is small. However, note that Assumption 3 is more likely to hold in the opposite case: when ρ is small and $|S|$ is large. Therefore, selecting ρ and $|S|$ requires balancing their effects on κ with ensuring that Assumption 3 holds.

Third, observe that the bound expression is invariant to ε provided $\varepsilon \in (0, 1)$. Intuitively, this occurs because the key object of interest in the proof is a sequence of Bernoulli variables (which contribute to the loss function error in a binary manner) that we use to bound the number of times the hidden action is misclassified by a constructed solution $\tilde{\pi}$. However, we note that ε does indeed appear in the non-dominant terms of the bound, as we would expect (see (EC.37) in the proof of Theorem 2).

Note that Theorem 2 is only valid for the case where each π_a is permitted to be any valid probability vector (i.e., $Q_\pi = \mathbb{R}^{m \times d}$). This additional condition is imposed on Π because the randomness of the set V can render the solution constructed by our proof approach infeasible for a more general parameter set Π . However, this additional assumption on Π is only needed for the finite-sample characterization of the error in Theorem 2; Proposition 2 below shows that the solution from the restricted estimator, $\tilde{\pi}_n$, is asymptotically optimal with respect to the exact estimator PA-C for any Π that satisfies (6).

PROPOSITION 2. *Let Assumption 3 hold. Then PA-D is asymptotically optimal with respect to PA-C: $|Z_n(\pi_n^*) - Z_n(\tilde{\pi}_n)| \rightarrow 0$.*

The asymptotic optimality established in Proposition 2 provides assurance that PA-D is a reasonable approximation to PA-C when n is large, which is precisely the regime where PA-C is likely to be intractable. As a consequence, we should also expect the restricted estimator to produce “good” estimates of π^0 for larger sample sizes. Having established that PA-D reasonably approximates PA-C, we now focus on developing a solution technique for tackling the mixed-integer program PA-D.

4.3. Statistical column generation

Observe that the size of the optimization problem PA-D grows with the number of candidate distributions in V , which can make PA-D computationally intractable if V is large. In this section, we propose a solution algorithm that involves solving PA-D over a subset of V – which we shall call V^+ – which dramatically improves the tractability of the estimator PA-D, with minimal degradation in estimation error. Because each candidate distribution in V is mapped to a set of decision variables in PA-D (where the set S indexes the distributions in V), our solution technique can be interpreted as a column generation algorithm.

The key step of our approach is a series of non-parametric hypothesis tests, which identifies a subset V^+ by performing pairwise comparisons of candidate distributions in V . The intuition is as follows. Consider any candidate distribution $\mathbf{v}_s \in V$, and recall from Algorithm 1 that \mathbf{v}_s is the empirical mass function over outcomes associated with the contracts in the s^{th} cluster. If there exists another cluster s' such that all contracts in clusters s and s' induce the same action from the agent, then \mathbf{v}_s and $\mathbf{v}_{s'}$ can be interpreted as two empirical mass functions that were generated by the same probability distribution (i.e., one of the rows of π^0). Therefore, our goal will be to apply non-parametric hypothesis tests to identify whether any pairs in V are generated by the same distribution, and to discard those that are effectively “duplicates”.

4.3.1. Hypothesis test function. A hypothesis test typically consists of four main steps: (1) a null hypothesis is specified that we wish to test, (2) a significance level α (i.e. Type I error rate) is specified for the test, (3) a test statistic is computed based on the sample data, and (4) the null hypothesis is rejected if and only if the magnitude of the test statistic exceeds a threshold τ_α , where τ_α depends on α . In the context of our column generation algorithm, the null hypothesis we will test is whether two candidate distributions \mathbf{v}_s and $\mathbf{v}_{s'}$ are generated from the same probability distribution (i.e., the same π_a^0), for many pairs (s, s') .

We first introduce some additional definitions that are required by our algorithm. For each $s \in S$, define a vector $\psi_s \in \mathbb{Z}_+^d$, where the j^{th} entry is the frequency of outcome j in the s^{th} cluster of Algorithm 1. The vector ψ_s is simply a convenient form for representing the candidate distribution \mathbf{v}_s within our hypothesis tests. Let $n_s = |I_s|$ be the number of observations in cluster s , and note

$n_s = \sum_{j \in J} \psi_{sj}$ for $s \in S$. Next, note that for each $s \in S$, by the weak law of large numbers there exists $\boldsymbol{\nu}_s \in \mathbb{R}_+^d$ such that $\|\boldsymbol{\nu}_s - \mathbf{v}_s\| \rightarrow 0$ as $n_s \rightarrow \infty$. We now define the main ingredient of the algorithm, which is a *test function* that declares whether $\boldsymbol{\psi}_s$ and $\boldsymbol{\psi}_{s'}$ are statistically different at a significance level α .

DEFINITION 1. $H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) : \mathbb{Z}_+^d \times \mathbb{Z}_+^d \mapsto \mathbb{R}$ is a test function if $\Pr(H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) > 0 | \boldsymbol{\nu}_s \neq \boldsymbol{\nu}_{s'}) \rightarrow 1$ as $n_s \rightarrow \infty$ and $n_{s'} \rightarrow \infty$ and $\Pr(H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) > 0 | \boldsymbol{\nu}_s = \boldsymbol{\nu}_{s'}) \leq \alpha$.

Definition 1 states that the hypothesis test function returns a positive value if and only if the null hypothesis – that the candidate distributions \mathbf{v}_s and $\mathbf{v}_{s'}$ are generated by the same probability distribution – is rejected. This definition subsumes many two-sample, non-parametric hypothesis tests. One example is the Kolmogorov–Smirnov hypothesis test (Massey 1951, Stephens 1974), which is widely used for its ease of implementation. In particular, the test function is given by

$$H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) = \sup_{j \in J} \left| \frac{\psi_{sj}}{n_s} - \frac{\psi_{s'j}}{n_{s'}} \right| - \tau_\alpha \sqrt{\frac{n_s + n_{s'}}{n_s n_{s'}}},$$

where τ_α^{KS} is the critical value associated with a significance level of α (Smirnov 1948). We note here that the Kolmogorov–Smirnov test is known to be conservative for discrete distributions (Slakter (1965), Conover (1972)). As a result, selecting τ_α based on Kolmogorov–Smirnov critical values for continuous distributions makes α an upper bound on the true Type I error rate in our setting, but otherwise does not affect the validity of our algorithm. Other examples of non-parametric tests that fit within our framework are the Anderson–Darling (Anderson et al. 1952, Scholz and Stephens 1987), Chi-squared (Cochran 1952) and the Cramér–von Mises (Anderson 1962) tests.

4.3.2. Algorithm overview. Let S^+ index the candidate distributions in V^+ . We let PA-D(S^+) denote formulation PA-D where S is replaced with the subset S^+ , and we let PA-D(S) denote the original formulation with the full set V . Let $V^- = V \setminus V^+$ and $S^- = S \setminus S^+$ denote the omitted distributions and the accompanying index set, respectively. Given a significance level α , we shall say two candidate distributions \mathbf{v}_s and $\mathbf{v}_{s'}$ are *statistically different* if and only if $H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) > 0$; that is, the null hypothesis that \mathbf{v}_s and $\mathbf{v}_{s'}$ were generated from a common probability distribution is rejected. In each iteration of the main loop of the algorithm, we perform a series of hypothesis tests identify a new candidate distribution to be introduced to V^+ , and solve PA-D(S^+) once there does not exist any distribution in V^- that is statistically different from every distribution in V^+ at a significance level of α . Specifically, in each iteration we compute

$$s^* = \operatorname{argmax}_{s \in S^-} \inf_{s' \in S^+} H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}).$$

Intuitively, \mathbf{v}_{s^*} is the distribution in V^- that is the “most” different from all distributions in V^+ , based on the selected test function H_α . The distribution \mathbf{v}_{s^*} is then added to V^+ if and only if

$$\inf_{s' \in S^+} H_\alpha(\boldsymbol{\psi}_{s^*}, \boldsymbol{\psi}_{s'}) > 0. \quad (15)$$

If (15) holds, then \mathbf{v}_{s^*} is statistically different from every distribution in V^+ , and is thus added to V^+ . If (15) does not hold, then there are no remaining distributions in V^- that are statistically different from all distributions in V^+ . In this case, we solve PA-D(S^+), and the algorithm terminates. A summary is given in Algorithm 2.

Algorithm 2: Statistical column generation (PA-D+)

Input: Data (\mathbf{r}^i, ξ^i) , $i \in I$, candidate distributions V produced by Algorithm 1, significance level $\alpha > 0$.

Initialize: Set $t = 0$. Select any $s \in S$. Set $S^+ = \{s\}$ and $S^- = S \setminus \{s\}$.

1. Let $s^* = \operatorname{argmax}_{s \in S^-} \inf_{s' \in S^+} H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'})$.

if $\inf_{s' \in S^+} H_\alpha(\boldsymbol{\psi}_{s^*}, \boldsymbol{\psi}_{s'}) \leq 0$ or $S^- = \emptyset$,

 Solve PA-D(S^+) and obtain solution $\boldsymbol{\pi}_n^+$, set $T = t$, and **terminate**.

else Update $t \leftarrow t + 1$, $S^+ \leftarrow \{S^+, s^*\}$, and $S^- \leftarrow S^- \setminus \{s^*\}$. Return to Step 1.

Output: Parameter estimate $\boldsymbol{\pi}_n^+$, iteration count T .

We will use “PA-D+” to denote the estimator represented by Algorithm 2. There are two main differences between existing column generation methods for large-scale integer programs and the one we propose in Algorithm 2. First, the column generation process in Algorithm 2 involves performing several hypothesis tests, which are fast to compute. By comparison, existing methods for integer programs typically generate columns by solving an auxiliary optimization problem (often called the *pricing* problem due to its use of dual information), which is often an integer program itself and may be difficult to solve (Lubbecke and Desrosiers 2005). Second, Algorithm 2 is not guaranteed to produce an optimal solution to PA-D; in contrast, the purpose of existing column generation methods is to solve the “original” optimization problem exactly. Therefore, Algorithm 2 effectively sacrifices optimality for computational efficiency. However, although Algorithm 2 does not produce optimal solutions to PA-D, it can be shown to produce a consistent estimate of $\boldsymbol{\pi}^0$, which is our main objective in this paper.

4.3.3. Consistency and iteration bound. Next, we present the main result of §4: Theorem 3 shows that the approximate solution obtained by Algorithm 2 preserves the consistency of the exact estimator PA-C.

THEOREM 3. *Let $\boldsymbol{\pi}_n^+$ be the estimate obtained by PA-D+ (Algorithm 2). Then*

$$\boldsymbol{\pi}_n^+ \longrightarrow \boldsymbol{\pi}^0.$$

As a consequence of Theorem 3, we should expect π_n^+ to provide a reasonable estimate of the unknown parameter π^0 . However, note that Theorem 3 is an asymptotic result only, and that for small n the estimate from PA-D+ may be less accurate than the exact estimate obtained by solving PA-C. We compare the performance of these two approaches numerically in §4.4.

Note that because the termination condition in Algorithm 2 depends on the outcome of a series of hypothesis tests, the total number of iterations, denoted by T , is a random variable. In Theorem 4 below, we show that $\mathbb{E}[T]$ is bounded by a function of the problem parameters, including the significance level α used in the hypothesis testing step of Algorithm 2.

THEOREM 4. *Let Assumption 3 hold. Further, assume that for each $s \in S$, there exists $a \in A$ such that $B_s \subseteq R_a(\pi^0)$. Then*

$$\mathbb{E}[T] \leq m[1 + \alpha \cdot |S| \cdot (|S| - m)].$$

The proof of Theorem 4 relies on upper bounding $\Pr(T > m)$ – the probability that the number of iterations in Algorithm 2 exceeds the number of agent actions. In particular, we show in the proof of Theorem 4 that $\Pr(T > m) \leq \alpha m S$. The intuition for the preceding inequality is as follows. Observe that by construction, the candidate distribution v_s is the empirical distribution over outcomes associated with all contracts \mathbf{r}^i such that $\mathbf{r}^i \in B_s$. Because for each $s \in S$, $B_s \subseteq R_a(\pi^0)$ for some $a \in A$ (by assumption), there are at most m unique distributions from which the empirical distributions v_s are generated, which are π_a^0 , $a \in A$. Next, note that in Algorithm 2, a candidate distribution is only added to the set V^+ if the hypothesis testing step finds it to be statistically different from every distribution in V^+ . Therefore, the event $\{T > m\}$ implies that a Type I error has occurred at some point during Algorithm 2; that is, a candidate distribution was added to V^+ despite the underlying distribution π_a^0 already being “represented” in V^+ by another candidate distribution.

Because α bounds the probability of making a Type I error, smaller values of α will make Algorithm 2 more *conservative* in adding new distributions to V^+ , thus increasing the probability of the event $\{T > m\}$. Conversely, if α is large, then it becomes more likely that a given distribution v_s is determined to be statistically different from those in V^+ , which leads to more distributions being added to V^+ , and thus a greater number of iterations. The dependence on S arises for a similar reason; as S increases, so does the number of omitted distributions V^- , which increases the likelihood that there exists a distribution in V^- that satisfies the inclusion criterion in Step 2 of Algorithm 2.

Additionally, note that the bound $\mathbb{E}[T] \leq |S|$ holds trivially, because $T = |S|$ implies $S_T^- = \emptyset$ by Algorithm 2. As a result, the bound in Theorem 4 may be vacuous if α is large, but is made meaningful for an appropriate selection of S and α . It is also straightforward to verify that the assumption in the statement of Theorem 4 implies that $|S| \geq m$, which confirms that the bound on $\mathbb{E}[T]$ is strictly positive for all $\alpha > 0$.

4.4. Numerical performance

In this section, we compare the performance of three estimation methods using synthetic data. The first two are solving the exact estimator (PA-C) and the restricted estimator (PA-D) directly with optimization software. The third is solving the restricted estimator using the column generation technique outlined in Algorithm 2 (PA-D+). We focus our comparison on the solution times and estimation errors from the three approaches.

4.4.1. Setup. Recall that m and d denote the number of actions and outcomes, respectively. We consider five problem sizes, given by $(m, d) \in \{(2, 2), (4, 5), (5, 10), (10, 20), (20, 40)\}$. For each of the five problem sizes, we consider three sample sizes, given by $n \in \{100, 500, 1000\}$. Then for each combination (m, d, n) , we randomly generate $\boldsymbol{\pi}^0$ from the appropriately-sized parameter set Π given by (6), where Q_π is given by (7). For each (m, d, n) , we randomly generate contract data by sampling \mathbf{r}^i uniformly from $[1, 10]^d$ for each $i = 1, \dots, n$, and sampling \mathbf{c} uniformly from $[0, 1]^m$. The outcome associated with each \mathbf{r}^i is obtained by solving the agent’s problem (1) under the corresponding $\boldsymbol{\pi}^0$. We repeat this procedure for a total of 10 trials for each (m, d, n) . To parameterize PA-D, we set $S = 50$ and $\rho = 10 \times d$. For the hypothesis testing step for PA-D+, we use the discrete analogue of the two-sample Anderson–Darling test (Scholz and Stephens 1987), and set $S = 50$, $\rho = 10 \times d$ and $\alpha = 0.05$. We use the optimization solver Gurobi 8.0 to solve PA-C, PA-D and PA-D+.

4.4.2. Results. Table 1 summarizes the average solution time and estimation errors over 10 trials for the three estimators. In each trial, the error associated with PA-C, PA-D and PA-D+ is given by $\frac{1}{md}\|\boldsymbol{\pi}^0 - \boldsymbol{\pi}_n^*\|$, $\frac{1}{md}\|\boldsymbol{\pi}^0 - \tilde{\boldsymbol{\pi}}_n\|$, and $\frac{1}{md}\|\boldsymbol{\pi}^0 - \boldsymbol{\pi}_n^+\|$, respectively. In all trials, we set a time limit of 3600 CPU seconds. Dashes in the table indicate instances where an optimal solution was not found within 3600 CPU seconds for any of the 10 trials. In many of these trials, no feasible solution was found within 3600 CPU seconds; we therefore only include errors obtained at optimal solutions to PA-C or PA-D when reporting the average estimation error.

We offer a few observations regarding Table 1. First, note that for each problem size, the estimation error generally decreases in n , which corroborates our consistency results (Proposition 1 and Theorem 3, respectively). Second, for smaller problem instances (e.g., $m = 4$, $d = 5$, $n = 1000$), PA-C is less computationally expensive than PA-D+, which we posit is a consequence of requiring fewer binary decision variables. However, PA-D+ generally scales more efficiently in the problem and sample size than PA-C and PA-D, with the most notable performance improvement occurring at larger problem instances (e.g., $m = 10$, $d = 20$, $n = 1000$). Third, solving the restricted estimator PA-D directly with Gurobi is less tractable than solving the exact estimator PA-C with Gurobi. This is again likely attributable to PA-D requiring many more binary variables than PA-C, due

to how the restricted parameter set is represented in the formulation PA-D. Nonetheless, the results indicate that this intractability can be overcome by (approximately) solving the restricted estimator using the statistical column generation technique, without significantly compromising estimation error. Fourth, observe that larger problem sizes are not necessarily more computationally expensive – for example, the average solution time of the instances $(2, 2, 1000)$ and $(5, 10, 1000)$ for PA-C was 245 and 12 seconds, respectively. We conjecture that this is because the larger problem sizes offer the estimator additional degrees of freedom in fitting the agent model to the data (due to containing a larger number of unknown parameters), which allows the optimization problem to more quickly attain the minimal objective function value. Lastly, the favorable performance of PA-D+ in the larger instances (e.g., $m = 20, d = 40$) suggests that our estimator and algorithm can also be used to tractably approximate contracting problems with continuous actions and outcomes through discretization.

Note that the purpose of Algorithm 2 is not to generate a provably optimal solution to PA-D, which is typically the case with similar column generation methods. Instead, our primary goal is to generate an estimate of the true parameter π^0 that is statistically consistent, competitive with solutions from solving the exact estimator, and attainable in a computationally efficient manner. Theorem 3 and the numerical results in Table 1 suggest that Algorithm 2 meets each of these criteria.

m	d	n	PA-C		PA-D		PA-D+	
			Time	Error	Time	Error	Time	Error
2	2	100	2	0.07	20	0.06	2	0.09
2	2	500	19	0.06	3432	0.06	4	0.06
2	2	1000	245	0.06	–	–	15	0.06
4	5	100	0	0.05	–	–	4	0.09
4	5	500	2	0.05	–	–	18	0.06
4	5	1000	3	0.05	–	–	66	0.06
5	10	100	1	0.04	–	–	4	0.06
5	10	500	6	0.03	–	–	14	0.04
5	10	1000	12	0.03	–	–	47	0.03
10	20	100	2404	0.02	–	–	3	0.02
10	20	500	–	–	–	–	15	0.01
10	20	1000	–	–	–	–	26	0.01
20	40	100	–	–	–	–	2	0.01
20	40	500	–	–	–	–	84	0.01
20	40	1000	–	–	–	–	211	0.01

TABLE 1. Solution time (CPU seconds) and normalized estimation error of three formulations averaged over 10 trials. Instances that did not solve to optimality under 3600 CPU seconds are omitted when calculating average estimation error. Dashes indicate no instance solved to optimality within 3600 CPU seconds in any trial.

5. Empirical Study: Randomizing Incentives in a Crowdwork Platform

In this section, we demonstrate our method by using it to investigate the effect of financial incentives on work quality in an online labor platform. First, we conducted an experiment on a crowdwork platform (Amazon Mechanical Turk) by randomly assigning incentive contracts to a pool of workers completing the same task. We then estimate an agent model from the experimental data, which allows us to characterize the link between incentives and quality and solve for an optimal incentive contract.

5.1. Background: Incentives and quality on Amazon Mechanical Turk

Crowdwork platforms are used by businesses that require temporary labor to complete tasks that are typically difficult for computers but simple for humans. Common tasks include audio transcription, classification of images, and data entry. The largest and most well-known crowdwork platform is Amazon’s Mechanical Turk (“mTurk”), which has been estimated to have 100,000 unique workers, with 2,000 active at any given time (Difallah et al. 2018).

The mTurk platform allows “requesters” to post tasks, along with a reward to be paid to the worker upon successful completion. Workers can select the tasks they want to complete, typically on a first-come, first-served basis. Requesters have discretion over whether to pay workers for their submissions, and can deny payment if the worker’s submission is incomplete or low quality.

Requesters can also provide bonuses to workers. Workers can be informed of the structure of the bonus payment within the instructions for a task, which offers the requester considerable flexibility in designing incentives.

The question of whether financial incentives improve quality of work in crowdwork platforms has been addressed in multiple studies, with differing conclusions. Mason and Watts (2009) find that incentives improve the quantity, but not quality of work; similarly, Yin et al. (2013) find that the magnitude of the bonus does not affect quality. In contrast, Horton and Chilton (2010) and Harris (2011) both find that quality can improve with worker pay. An important study in this line of research is by Ho et al. (2015), who suggest that for tasks where quality plausibly depends on worker effort (e.g., proofreading), incentives can improve quality.

With respect to experimental design, we underline two differences between our study and the work cited above. First, instead of assigning workers to a finite number of treatments (e.g., bonus or no bonus), we vary incentives in a continuous manner, meaning the parameters of the incentive contract are randomly drawn *for each worker*. This design significantly complicates the implementation of the experiment on mTurk, but introduces useful variation for estimating our agent model. Second, we examine how incentives affect the *distribution* of work quality, instead of average quality.

5.2. Experimental setup

5.2.1. Task design. A major source of observable heterogeneity in the mTurk worker population is location. Approximately 91% of workers are located in two countries: the US (75%) and India (16%) (Difallah et al. 2018). We collected and analyzed data from both countries separately.

The experiment involved posting two types of tasks on mTurk. First, we posted a recruitment task in which workers were paid \$1.00 for agreeing to be notified of future tasks by email. We recruited 250 workers from both the US and India using this task, for a total worker pool of 500. The recruitment task in each country was made available for one day, and reached its maximum number of submissions (250 for each location) within 3 hours of posting. Second, inspired by Ho et al. (2015), we created a proofreading task by inserting 10 typos into a one-page, 500-word excerpt from a newspaper article. The proofreading task required workers to report the line number and correct spelling for each misspelled word in the article (e.g., “5:automobile”). We use a proofreading task because it allows us to objectively measure the quality of each submission (percentage of typos identified). After constructing the worker pool, we posted the proofreading task on mTurk and notified each worker by email of the task’s availability. The task was available for 24 hours.

5.2.2. Incentive structure and randomization. We next describe how we randomized incentives among workers. The mTurk platform allows requesters to assign “qualification” criteria to tasks, which only allows workers with the required qualifications to view and complete the task. For example, a requester might assign a location or age-based qualification to a task if they wish to target a specific worker population. Requesters can also create and assign custom qualifications to workers. When conducting a randomized experiment, creating and randomly assigning qualifications to workers effectively allows the requester to construct multiple treatment groups, where each qualification represents one treatment.

We use the qualification feature in mTurk to create exogenous variation in worker incentives. We first created 500 unique qualifications, and assigned each qualification to a single worker in the pool. We then created 500 tasks where each task was randomly assigned to a qualification. As a result, for each of the 500 tasks, only a single worker in our pool was able to view and complete it.

The payment for the proofreading task consisted of two components: a *base* payment for finding at least 25% of the typos in the document, and an additional *bonus* payment for finding at least 75% of typos. For each task (i.e., each worker in the pool), we drew *base* and *bonus* uniformly from the interval [\$0.10, \$1.00], rounded to the nearest \$0.01. We provided the details of the payment structure upfront in the task instructions. Because workers were only able to view the task assigned to them, workers could not observe the payment offered to others, and had no knowledge that payments were randomized. Note that in the context of the proofreading task, worker output corresponds to the fraction of typos corrected, which we also refer to as the task *quality*.

5.2.3. Submissions. We collected a total of 346 submissions, each from a unique mTurk worker. Of these, 215 submissions were from US-based workers, and 131 were from India-based workers. We analyze the data from US and India workers separately throughout our study. Figure 1 depicts the distribution of quality scores for workers in each location. Note that a large number of submissions achieve a quality score of 0. Low-quality submissions are a well-known feature of mTurk; because verifying responses manually for a large number of submissions is difficult, workers may submit blank or low-quality responses in the hope of nevertheless receiving a payment (Ipeirotis et al. 2010). Scores of 0 may also be due to submissions not being in the correct format, which we specified as a condition for payment in the task instructions.

The mTurk platform provides timestamps for when a worker accepted and submitted a task. The mean completion time (i.e., time between acceptance and submission) was 9.7 minutes, and 95% of workers submitted the task between 1 and 29 minutes after accepting it. Because mTurk allows workers to accept tasks into a queue before working on them, the recorded completion time is an upward-biased measurement of the actual time the worker spent on the task. As a result,




Figure 1 Distribution of quality scores for submissions made by workers in the US and India.

completion time may be a poor proxy for true worker effort, because the requester cannot observe how much time the task spent in the worker’s queue. We therefore treat effort as fully hidden, and do not use completion time data in our study.

Based on each worker’s completion time, we estimated the average wage to be \$14.50/hr for our task (including the guaranteed \$1.00 payment at recruitment). This is likely a conservative estimate of the true average wage due to the queuing behavior described above.

5.3. Estimation and validation

Next, we describe the application of our estimation procedure to the experimental data. Putting the results of the experiment in the format required by our estimator is straightforward. Recall that each worker was eligible for three possible payments based on their submission quality: no payment (if they found 0-25% of typos), a base payment (25%-75%), or both a base and bonus payment (75-100%). In our framework, this corresponds to $d = 3$ possible performance levels for the worker’s outcome ξ^i . Accordingly, the i^{th} worker’s incentive contract \mathbf{r}^i has the components $r_1^i = 0$, $r_2^i = \text{base}^i$, and $r_3^i = \text{base}^i + \text{bonus}^i$, where base^i and bonus^i are the randomly generated parameters for that worker. For the PA-D+ algorithm, we set $\rho = 0.5$, $S = 10$ and $\alpha = 0.0001$ throughout all experiments.

5.3.1. Measuring goodness of fit. We require a goodness of fit metric for fitting and validating the model. Recall that our estimation procedure generates a prediction of the outcome distribution: Given an estimate $\hat{\pi}$, a contract \mathbf{r} , and action costs \mathbf{c} , the model’s prediction of the outcome distribution under \mathbf{r} is $\hat{\pi}_{a(\mathbf{r})}$, where $a(\mathbf{r})$ is the agent’s optimal action under contract \mathbf{r} .

For ease of interpretation, we measure goodness of fit as the absolute error between the empirical and predicted probability of a given outcome, averaged over all outcomes. Specifically, let (\mathbf{r}^i, ξ^i) , $i = 1, \dots, n$ be the data we wish to measure our model fit against. As before, for each i , let $y_j^i = 1$ if $\xi^i = j$ (if outcome j is observed). Then the mean absolute error (MAE) is given by

$$\text{MAE} = \frac{1}{d} \sum_{j=1}^d \left| \frac{1}{n} \sum_{i=1}^n (\hat{\pi}_{a(\mathbf{r}^i), j} - y_j^i) \right|. \quad (16)$$

5.3.2. Setting cost parameters. Two hyperparameters in our model are the number of actions, m , and the action costs, c_1, \dots, c_m . We selected these parameters using a standard 10-fold cross-validation procedure, using MAE to measure cross-validation errors. To avoid performing an extremely large number of cross-validation iterations, we imposed additional structure by assuming action costs were of the form $c_a = (a - 1) \cdot \delta$, for $a = 1, \dots, m$. We used cross-validation to jointly select m and δ from the sets $\{2, 3, 4\}$ and $\{0.02, 0.05, 0.1, 0.2, 0.5\}$, respectively (units of the latter set are dollars). Results are presented in Table 2. Note that errors are relatively stable for all values of δ when $m = 2$ or $m = 3$, whereas the model appears to overfit for $m = 4$. We select $(m, \delta) = (3, 0.1)$ for both the US and India datasets, resulting in the cost vector $\mathbf{c} = [0, 0.1, 0.2]$. Lastly, for Algorithms 1 and 2, we set $\rho = 0.5$, $S = 50$, and $\alpha = 0.01$, and use a Chi-squared test for the hypothesis test step of Algorithm 2.

We note that our handling of action costs is fairly stylized, because they are treated as hyperparameters to be tuned prior to model fitting. Horton and Chilton (2010) estimates the median reservation wage of mTurk workers to be \$1.38/hr. Given that the median completion time for our task was 8 minutes, \$0.00–\$0.20 appears to be a reasonable approximation for the range of effort costs of an mTurk worker. We discuss costs in more detail in §5.5.

5.3.3. Bootstrapping. Given our moderately-sized data set ($n = 215$ and $n = 131$), we validated our estimation procedure by bootstrapping. For each of 100 repetitions, we sampled n

		δ				
		0.02	0.05	0.1	0.2	0.5
US	2	0.06	0.06	0.07	0.06	0.06
	3	0.04	0.06	0.04	0.06	0.05
	4	0.18	0.11	0.12	0.12	0.08
IN	2	0.07	0.08	0.06	0.06	0.08
	3	0.07	0.05	0.06	0.06	0.08
	4	0.15	0.08	0.06	0.11	0.09

TABLE 2. 10-fold cross-validation errors (MAE) for US and India groups, with varying number of actions (m) and cost spacing (δ).

observations with replacement, and estimated the model parameters from the sample using Algorithm 2. For each repetition, we assessed model fit using two hypothesis tests: a Chi-squared (χ^2) test, which is appropriate in our setting because outcomes are discrete, and an exact test using MAE as the test-statistic, where the sampling distribution is obtained through Monte Carlo simulation. In both hypothesis tests, the null hypothesis is that the empirical distribution of quality outcomes in the out-of-bootstrap data is generated by the fitted model. Accordingly, we interpret large p -values as indicating a good model fit.

Table 3 shows the distribution of test statistics and associated p -values over the 100 bootstrap repetitions. Both tests produced comparable p -values within each worker group. Note that the median p -value was above 0.1 for both groups, which suggests the model reasonably fits the joint distribution over (\mathbf{r}, ξ) in the majority of bootstrap iterations.

		5 th	25 th	Median	75 th	95 th
US	χ^2 (p -value)	0.20 (0.90)	1.70 (0.43)	3.42 (0.18)	7.36 (0.03)	23.58 (0.00)
	MAE (p -value)	0.02 (0.89)	0.04 (0.42)	0.07 (0.14)	0.10 (0.02)	0.15 (0.00)
IN	χ^2 (p -value)	0.18 (0.91)	0.99 (0.61)	2.13 (0.34)	5.23 (0.07)	14.33 (0.00)
	MAE (p -value)	0.02 (0.93)	0.05 (0.65)	0.09 (0.35)	0.14 (0.09)	0.19 (0.00)

TABLE 3. Percentiles of Chi-squared (ξ^2) and MAE test statistics with associated p -values over 100 bootstrap repetitions.

Table 4 presents the estimated values of $\boldsymbol{\pi}$ and standard errors for both worker groups. Each 3×3 section in the center of Table 4 corresponds to the estimated $\boldsymbol{\pi}$ matrix for the labelled worker group, averaged over 100 bootstrap repetitions. For convenience, we refer to the outcome in which the worker earns the bonus ($\xi^i = 3$) as the “bonus outcome”, and the probability that this outcome is realized as the “bonus probability”. Note that the highest cost action ($a = 3$) has the highest bonus probability in both worker groups, and that the bonus probability is lower in the India worker group compared to the US group, for all actions.

Our estimation procedure treats each action a as a latent variable. The solution to the estimation problem produces a clustering where each outcome is assumed to have been generated by one of the m distributions (i.e., agent actions). As a result, for each bootstrap repetition, we can count the number of observations that are assigned to each action by the estimator. The average number of observations mapped to each action are reported in the final column of Table 4.

		Outcomes (j)			# Obs.
		1	2	3	
US	1	0.46 (0.18)	0.34 (0.18)	0.20 (0.09)	19
	2	0.30 (0.10)	0.42 (0.12)	0.28 (0.08)	27
	3	0.20 (0.07)	0.43 (0.07)	0.37 (0.06)	169
IN	1	0.57 (0.12)	0.34 (0.12)	0.09 (0.07)	36
	2	0.45 (0.09)	0.38 (0.11)	0.17 (0.06)	36
	3	0.35 (0.07)	0.42 (0.07)	0.23 (0.06)	59

TABLE 4. Estimated values of π for both groups, with standard errors in parentheses. The final column reports the number of in-bootstrap observations mapped to each action, averaged over 100 bootstrap repetitions.

5.3.4. Predictive performance. Next, we evaluate the predictive performance of the estimator. For each of the 100 bootstrap models, we compute the prediction error (given in (16)) attained by the fitted model on the out-of-bootstrap observations. We set $S = 10$, $\rho = 0.5$, and $\alpha = 0.0001$. To serve as performance benchmarks, we repeat the bootstrap procedure for standard implementations of multinomial logistic regression (MLR) and classification trees (CT), both of which also generate predictions of the outcome distribution for a given set of contracts.⁷ Figure 2 depicts the distribution over prediction errors for the three methods over the 100 bootstrap repetitions. For the US data, the average MAE for PA-D+, MLR and CT is 0.059, 0.070 and 0.093, respectively; for the India data, the average errors are 0.072, 0.086, and 0.112. In summary, Figure 2 confirms that the PA-D+ estimator produces sound predictions on the experimental mTurk data, and is competitive with well-known benchmark methods. In §EC.2 of the electronic companion, we further compare all three methods on several synthetic instances, and find that our estimator continues to perform well.

⁷ Both benchmark methods are implemented using MATLAB’s Statistics and Machine Learning Toolbox using default settings.





Figure 2 Comparison of out-of-bootstrap prediction errors for PA-D+, multinomial logistic regression (MLR) and classification trees (CT) on mTurk data (100 repetitions).

5.4. Impact of bonuses on quality


We now use the estimated model to examine the effect of varying the bonus payment on quality. For a given incentive contract, we form a prediction of the outcome distribution by averaging over the 100 bootstrapped models, which improves stability and reduces overfitting (Breiman 1996). Let $\hat{\pi}^1, \dots, \hat{\pi}^K$ be the estimates obtained from K bootstrap repetitions. The probability of observing outcome j under the incentive contract \mathbf{r} is then given by $\frac{1}{K} \sum_{k=1}^K \hat{\pi}_{a^k(\mathbf{r}),j}^k$, where $a^k(\mathbf{r})$ is the optimal action for contract \mathbf{r} in the k^{th} agent model. To isolate the influence of the bonus payment, we fix the base payment to \$0.10, vary the bonus payment between \$0.10 and \$1.00, and compute the probability of each quality outcome under each bonus amount. We repeat for a base payment of \$1.00.

Figure 3 shows the results for both the US and India groups of workers. For a base payment of \$0.10 (Figure 3a and 3b), the bonus probability (i.e., probability that submission quality is above 75%) increases moderately for both groups as the bonus is increased from \$0.10 to 1.00 (from 0.21 to 0.36 for the US group; from 0.09 to 0.17 for the India group). However, with a base payment of \$1.00 (Figure 3c and 3d), the effect of increasing the bonus payment from \$0.10 to \$1.00 is dampened (bonus probability increases from 0.34 to 0.37 for the US group; 0.18 to 0.22 for the India group). These results suggest that increasing the bonus payment can indeed increase quality, but the effect is significantly diminished when the base payment is already high. A qualitatively similar result can be obtained by fixing the bonus payment and varying the base payment (results not shown).

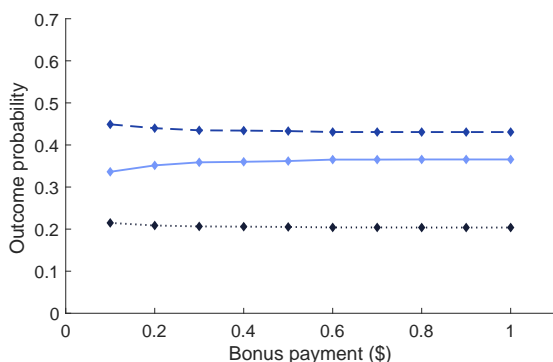
We shed some light on the mechanics behind Figure 3. Because our predictions are based on the average of 100 different agent models, for a fixed incentive contract, we can count the number of models in which each action is taken. Further, note that if the bonus payment increases, an agent may find it optimal to “switch” from a low-cost action to a high-cost action, thus increasing the




(a) US, base = \$0.10.



(b) India, base = \$0.10.



(c) US, base = \$1.00.



(d) India, base = \$1.00.

Figure 3 Effect of varying bonus payment on probability of each quality outcome (0-25%, 25-75%, 75-100%) for US and India workers.

probability of realizing the high-quality outcome. The change in probabilities depicted in Figure 3 is the result of the underlying agent models jumping from one action to the next as the parameters of the contract change.

Figure 4 shows the fraction of agent models that take each of the three actions as the bonus is increased from \$0.10 to \$1.00. Note that the four panels in Figure 4 map to the four panels in Figure 3. As expected, when the base payment is \$0.10, increasing the bonus amount from \$0.10 to \$1.00 is associated with agents switching away from the lowest cost action ($a = 1$) toward the higher cost actions ($a = 2$ and $a = 3$). Moreover, the shift toward higher cost actions is more pronounced for the US worker group, where the fraction of agents taking the highest cost action ($a = 3$) increases from 0 to 0.69; for the India group, this fraction increases from 0 to 0.17. In parallel with Figure 3, when the base payment is \$1.00, the fraction of agents taking the highest cost action ($a = 3$) is higher overall, but the shift toward higher cost actions as the bonus is increased is muted. In other words, the stability in selected actions shown in Figures 4c and 4d explains the stability in outcome probabilities seen in Figures 3c and 3d. We emphasize here that Figure 4 is intended to

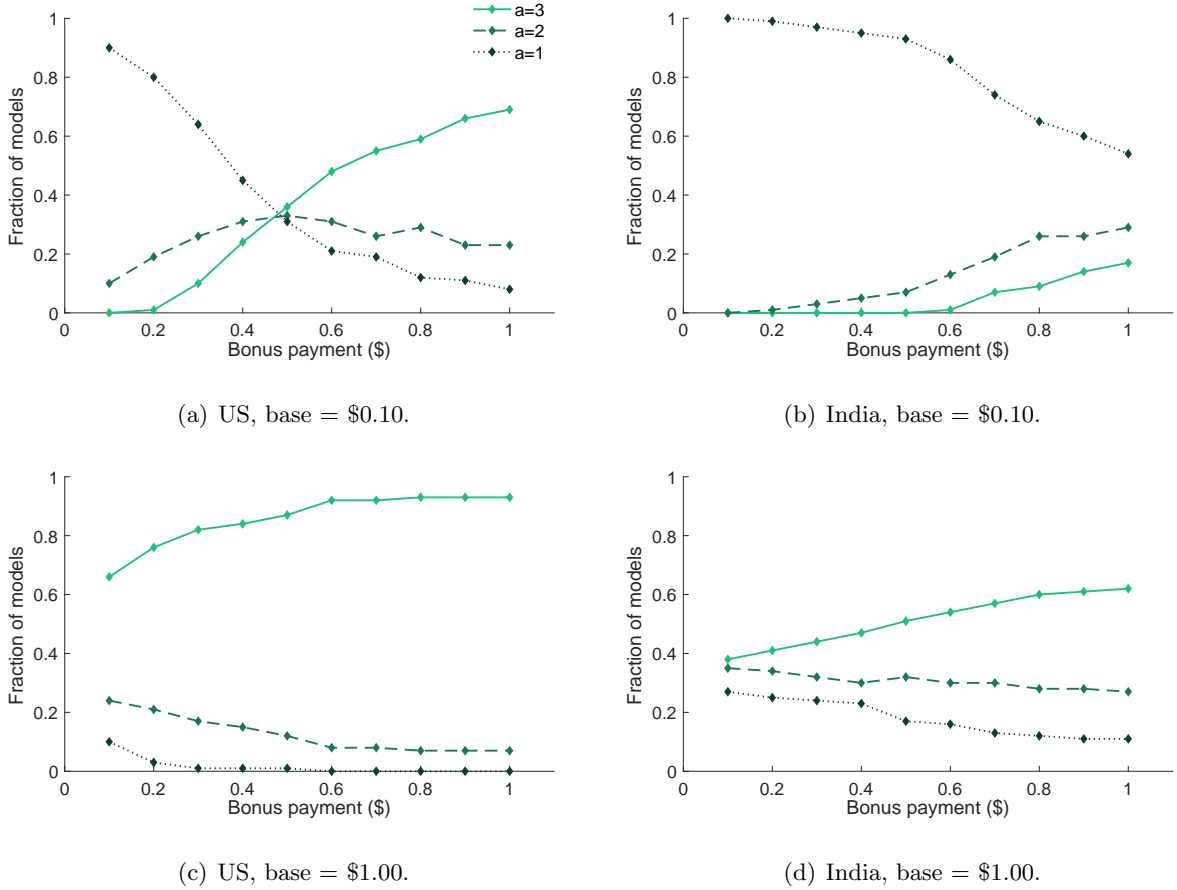


Figure 4 Effect of bonus payment on optimal agent actions in 100 bootstrapped models.

illustrate the mechanics behind the predictions in Figure 3, and is not necessarily a depiction of worker behavior.

5.5. Solving for an optimal incentive contract

An advantage of our model specification is that it leads to an optimal contracting problem that is highly tractable (see §EC.1 of the electronic companion for details). To illustrate this in the context of our mTurk study, we consider the simple problem of maximizing the bonus probability (i.e., outcome $\{\xi = 3\}$) subject to a budget constraint on the expected payment:

$$\text{maximize}_{\mathbf{r}} \hat{\pi}_{\hat{a}(\mathbf{r}),3} \quad (17a)$$

$$\text{subject to } \hat{a}(\mathbf{r}) = \underset{a \in A}{\operatorname{argmax}} \mathbf{r}^\top \hat{\boldsymbol{\pi}}_a - c_a, \quad (17b)$$

$$\mathbf{r}^\top \hat{\boldsymbol{\pi}}_{\hat{a}(\mathbf{r})} \leq \Gamma, \quad (17c)$$

$$\mathbf{r} \geq \mathbf{0}. \quad (17d)$$

The formulation above is a special case of the general optimal contracting problem presented in §EC.1 of the electronic companion, and can be solved exactly by solving $|A|$ linear programs. An




Figure 5 Frontier of optimal bonus probabilities under varying budget parameter Γ .

important consequence of the tractability of the optimal contracting problem (17) is that we can easily characterize the performance of the optimal contracts as the budget parameter Γ varies. To do so, we solve (17) for each $\Gamma \in \{0.05, 0.1, \dots, 1\}$ (for each of the 100 bootstrap estimates), and compute the average bonus probability under each value of Γ .

Figure 5 shows the resulting frontiers for both the US and India worker data. Because the curves are obtained by solving the optimal contracting problem (17), they represent estimates of the maximum attainable performance for both worker groups over the entire class of contracts used in the experiment. The value of the budget parameter Γ can be interpreted as the expected payment to the agent under the corresponding optimal contract. Our main finding is that higher payments increase quality modestly: increasing the expected payment from \$0.10 to \$1.00 increases the bonus probability under the optimal contract by 0.08–0.12, depending on the worker group. However, the most striking observation is that returns to quality diminish at fairly low payment levels, with quality improvements leveling off around \$0.30 and \$0.60 for the US and India groups, respectively (we discuss possible explanations in §5.7).

Figure 5 also clearly depicts the difference in the performance of optimal contracts between the US and India worker groups. For example, for the US group, attaining a bonus probability of 0.30 requires an expected payment of at least \$0.20; for the India group, a bonus probability of 0.30 is not attainable through higher payments alone. It can also be observed that the bonus probability is approximately 0.10–0.15 higher among US workers across all payment levels.

5.6. Experimental validation of contract performance

To validate the predicted performance of the optimal contracts shown in Figure 5, we conducted six follow-up experiments on mTurk. First, for each of the 100 bootstrap estimates $\hat{\pi}^1, \dots, \hat{\pi}^K$, we solved the optimal contracting problem (17) for $\Gamma \in \{0.25, 0.50, 0.75\}$, which corresponds to three different points on the frontiers in Figure 5. We then computed the optimal contract by taking the




Figure 6 Empirical bonus probabilities and 95% prediction intervals of six contracts implemented on mTurk.

component-wise average of the 100 solutions to (17). This produced six different testable contracts (i.e., combinations of the *base* and *bonus* parameters), which are shown in Table 5. We implemented each contract on mTurk by recruiting a new pool of 600 unique workers (using the same approach described in §5.2), and assigning 100 workers to each of the six contracts. Table 5 summarizes the results from these experiments, including the empirical bonus probability for each contract (i.e., the fraction of submissions with quality above 75%). In Figure 6, we plot the empirical bonus probabilities along with the 95% prediction intervals obtained from the bootstrap.

Figure 6 shows that for each of the six experimentally tested contracts, the empirical bonus probability sits comfortably inside its corresponding prediction interval, and is often close to the midpoint of the interval. In general, the prediction intervals are wide, which is unsurprising given that many other factors likely influence submission quality beyond the payment amount, including unobserved worker attributes. Further, validating the predictions from any model through experiments is challenging in general; because the worker population on mTurk is not temporally static (Difallah et al. 2018), the worker population in the validation experiments may be different from the initial experiments used to estimate the model. Nevertheless, our results in Figure 6 suggest that the estimator can reasonably predict experimental outcomes under a given incentive contract.

	Budget (Γ)	Base	Bonus	Submissions	> 75%	Empirical Bonus Probability
US	0.25	0.16	0.42	20	5	0.20
	0.50	0.23	0.47	52	17	0.33
	0.75	0.41	0.78	39	14	0.36
IN	0.25	0.11	0.31	73	9	0.12
	0.50	0.25	0.42	79	10	0.13
	0.75	0.58	0.77	71	8	0.11

TABLE 5. Optimal incentive contracts under three different values of Γ and associated results from mTurk experiments.

5.7. Discussion

Our results suggest larger incentives can increase quality on crowdwork platforms, corroborating the results of Ho et al. (2015). While similar results are reported in the literature, we have taken a complementary approach by characterizing worker performance over a *class* of incentive contracts. Further, the tractability of the optimal contracting problem under our agent model allows us to estimate performance under an optimal contract. In particular, as summarized in Figure 5, we find that increasing the expected worker payment by about \$1 increases the probability that a worker crosses the bonus threshold by 0.08–0.12, depending on the worker’s location. Most notably, we find diminishing returns to quality at relatively low payments in both worker groups, which may help explain why requesters tend to set low wages on mTurk (Hara et al. 2018).

We also observe that quality can depend strongly on the worker’s location. In particular, as seen in Figure 5, the bonus probability for the India group at an expected payment \$1.00 is comparable to the US group at \$0.10. This result aligns with a finding by Shaw et al. (2011), who observe that quality on mTurk is much more strongly associated with worker location than financial incentives. While we have only focused on worker location in this study, our approach can be readily extended to other worker attributes, provided sufficient data is available.

We highlight some limitations of our study and note directions for future work. First, we have treated agent costs as hyperparameters by tuning them through cross-validation. This makes the costs used in our model a rough approximation of actual worker costs, and may limit the interpretability of the resulting agent model. Our agent model also does not capture many of the worker dynamics present in crowdwork platforms. Horton and Chilton (2010) point out that mTurk worker output appears to deviate from what would be predicted by simple, rational agent models, which applies to our model as well. Lastly, an important aspect of crowdwork not addressed here is worker welfare. In particular, mTurk has been widely criticized for low worker pay, which is often far below the US minimum wage (Hara et al. 2018). While we did not address worker welfare in this paper, our modeling framework can also be used to characterize welfare over a class of incentive contracts,

and allows for welfare considerations to be explicitly incorporated into the optimal contracting problem (e.g., by imposing constraints on agent utility). Investigating the trade-off between worker welfare and quality in crowdwork may be a fruitful direction for future work.

6. Conclusion

We proposed an approach for estimating parameters that govern agent production in a moral-hazard principal-agent model. First, we presented an estimator for a non-parametric agent model, and showed it to be statistically consistent. To avoid computational drawbacks of solving the estimator exactly, we proposed an approximate estimator based on a restricted parameter set, and characterized the approximation error both asymptotically and in a finite-sample setting. To solve the restricted estimator, we developed a novel column generation technique that uses hypothesis testing to select variables, which we showed preserves consistency. Numerical results show that the approximation scheme and solution technique produce accurate estimates in a computationally efficient manner. Lastly, we applied our estimator to data from a randomized experiment on a crowdwork platform to demonstrate how our method can be used to characterize performance over a class of incentive contracts and identify optimal incentives from the estimated model.

We conclude by noting some possible directions for future work. Our estimation procedure is built upon a general but simple moral-hazard agent model; it may be useful to extend our approach to accommodate other common features of principal-agent models, such as unobserved heterogeneity and risk aversion. There may also be fertile ground in generalizing our statistical column generation algorithm to other integer programming problems. In particular, our approach may be relevant to other estimation problems where the parameter space is a very large set of discrete distributions. Lastly, estimating an agent model from data may be valuable for investigating questions related to worker welfare, which is an issue of increasing prominence in online labor platforms.

References

- Abreu, Dilip, David Pearce, Ennio Stacchetti. 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society* 1041–1063.
- Adams, Warren P, Richard J Forrester, Fred W Glover. 2004. Comparisons and enhancement strategies for linearizing mixed 0-1 quadratic programs. *Discrete Optimization* **1**(2) 99–120.
- Anderson, Theodore W. 1962. On the distribution of the two-sample Cramer-von Mises criterion. *The Annals of Mathematical Statistics* 1148–1159.
- Anderson, Theodore W, Donald A Darling, et al. 1952. Asymptotic theory of certain goodness of fit criteria based on stochastic processes. *The Annals of Mathematical Statistics* **23**(2) 193–212.
- Aswani, Anil, Zuo-Jun Shen, Auyon Siddiq. 2018. Inverse optimization with noisy data. *Operations Research* **66**(3) 870–892.

- Aswani, Anil, Zuo-Jun Max Shen, Auyon Siddiq. 2019. Data-driven incentive design in the Medicare Shared Savings Program. *Operations Research*.
- Bajari, Patrick, C Lanier Benkard, Jonathan Levin. 2007. Estimating dynamic models of imperfect competition. *Econometrica* **75**(5) 1331–1370.
- Barnhart, Cynthia, Ellis L Johnson, George L Nemhauser, Martin WP Savelsbergh, Pamela H Vance. 1998. Branch-and-price: Column generation for solving huge integer programs. *Operations Research* **46**(3) 316–329.
- Bertsimas, Dimitris, Vishal Gupta, Ioannis Ch Paschalidis. 2015. Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* **153**(2) 595–633.
- Bertsimas, Dimitris, John N Tsitsiklis. 1997. *Introduction to linear optimization*, vol. 6. Athena Scientific Belmont, MA.
- Bickel, Peter J, Kjell A Doksum. 2015. *Mathematical Statistics: Basic Ideas and Selected Topics, Volumes I-II Package*. Chapman and Hall.
- Breiman, Leo. 1996. Bagging predictors. *Machine learning* **24**(2) 123–140.
- Casella, George, Roger L Berger. 2002. *Statistical inference*, vol. 2. Duxbury Pacific Grove, CA.
- Chan, Timothy CY, Taewoo Lee, Daria Terekhov. 2018. Inverse optimization: Closed-form solutions, geometry, and goodness of fit. *Management Science*.
- Cochran, William G. 1952. The Chi-squared test of goodness of fit. *The Annals of Mathematical Statistics* 315–345.
- Conover, William J. 1972. A Kolmogorov goodness-of-fit test for discontinuous distributions. *Journal of the American Statistical Association* **67**(339) 591–596.
- de Zegher, Joann F, Dan A Iancu, Hau L Lee. 2017. Designing contracts and sourcing channels to create shared value. *Manufacturing & Service Operations Management*.
- Dietze, S, M Schäuble. 1985. On the relationship between fenchel and Lagrange duality for optimization problems in general spaces. *Optimization* **16**(1) 7–14.
- Difallah, Djellel, Elena Filatova, Panos Ipeirotis. 2018. Demographics and dynamics of mechanical turk workers. *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 135–143.
- Duffo, Esther, Rema Hanna, Stephen P Ryan. 2012. Incentives work: Getting teachers to come to school. *American Economic Review* **102**(4) 1241–78.
- Esfahani, Peyman Mohajerin, Soroosh Shafieezadeh-Abadeh, Grani A Hanasusanto, Daniel Kuhn. 2018. Data-driven inverse optimization with imperfect information. *Mathematical Programming* **167**(1) 191–234.

-
- Ferrall, Christopher, Bruce Shearer. 1999. Incentives and transactions costs within the firm: Estimating an agency model using payroll records. *The Review of Economic Studies* **66**(2) 309–338.
- Gayle, George-Levi, Robert A Miller. 2015. Identifying and testing models of managerial compensation. *The Review of Economic Studies* **82**(3) 1074–1118.
- Georgiadis, George, Michael Powell. 2021. A/B contracts .
- Glover, Fred. 1975. Improved linear integer programming formulations of nonlinear integer problems. *Management Science* **22**(4) 455–460.
- Grossman, Sanford J, Oliver D Hart. 1983. An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society* 7–45.
- Hara, Kotaro, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, Jeffrey P Bigham. 2018. A data-driven analysis of workers’ earnings on amazon mechanical turk. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- Harris, Christopher. 2011. You’re hired! an examination of crowdsourcing incentive models in human resource tasks. *Proceedings of the Workshop on Crowdsourcing for Search and Data Mining (CSDM) at the Fourth ACM International Conference on Web Search and Data Mining (WSDM)*. Hong Kong, China, 15–18.
- Ho, Chien-Ju, Aleksandrs Slivkins, Siddharth Suri, Jennifer Wortman Vaughan. 2015. Incentivizing high quality crowdwork. *Proceedings of the 24th International Conference on World Wide Web*. 419–429.
- Ho, Chien-Ju, Aleksandrs Slivkins, Jennifer Wortman Vaughan. 2016. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. *Journal of Artificial Intelligence Research* **55** 317–359.
- Holmstrom, Bengt. 1979. Moral hazard and observability. *The Bell Journal of Economics* 74–91.
- Horton, John Joseph, Lydia B Chilton. 2010. The labor economics of paid crowdsourcing. *Proceedings of the 11th ACM Conference on Electronic Commerce*. 209–218.
- Ipeirotis, Panagiotis G, Foster Provost, Jing Wang. 2010. Quality management on Amazon Mechanical Turk. *Proceedings of the ACM SIGKDD Workshop on Human Computation*. 64–67.
- Keshavarz, Arezou, Yang Wang, Stephen Boyd. 2011. Imputing a convex objective function. *2011 IEEE International Symposium on Intelligent Control*. IEEE, 613–619.
- Kuleshov, Volodymyr, Doina Precup. 2014. Algorithms for multi-armed bandit problems. *arXiv preprint arXiv:1402.6028* .
- Laffont, Jean-Jacques, David Martimort. 2009. *The theory of incentives: The principal-agent model*. Princeton University Press.
- Lee, Donald KK, Stefanos A Zenios. 2012. An evidence-based incentive system for Medicare’s End-Stage Renal Disease Program. *Management Science* **58**(6) 1092–1105.

- Lubbecke, Marco E, Jacques Desrosiers. 2005. Selected topics in column generation. *Operations Research* **53**(6) 1007–1023.
- Luedtke, James, Mahdi Namazifar, Jeff Linderoth. 2012. Some results on the strength of relaxations of multilinear functions. *Mathematical Programming* **136**(2) 325–351.
- Lyft. 2021. Ride Challenges. URL: <https://help.lyft.com/hc/enca/articles/360001943867-Ride-Challenges>.
- Mason, Winter, Duncan J Watts. 2009. Financial incentives and the performance of crowds. *Proceedings of the ACM SIGKDD Workshop on Human Computation*. 77–85.
- Massart, Pascal. 1990. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability* 1269–1283.
- Massey, Frank J. 1951. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association* **46**(253) 68–78.
- Misra, Sanjog, Anne T Coughlan, Chakravarthi Narasimhan. 2005. Salesforce compensation: An analytical and empirical examination of the agency theoretic approach. *Quantitative Marketing and Economics* **3**(1) 5–39.
- Misra, Sanjog, Harikesh S Nair. 2011. A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics* **9**(3) 211–257.
- Paarsch, Harry J, Bruce Shearer. 2000. Piece rates, fixed wages, and incentive effects: Statistical evidence from payroll records. *International Economic Review* **41**(1) 59–92.
- Postmates. 2021. How do bonuses and incentives work? URL: <https://support.postmates.com/fleet/articles/228603028-article-How-do-bonuses-and-incentives-work>.
- Radner, Roy. 1981. Monitoring cooperative agreements in a repeated principal-agent relationship. *Econometrica: Journal of the Econometric Society* 1127–1148.
- Rockafellar, R Tyrrell, Roger J-B Wets. 2009. *Variational analysis*, vol. 317. Springer Science & Business Media.
- Rogerson, William P. 1985. Repeated moral hazard. *Econometrica: Journal of the Econometric Society* 69–76.
- Sappington, David EM. 1991. Incentives in principal-agent relationships. *Journal of economic Perspectives* **5**(2) 45–66.
- Scholz, Fritz W, Michael A Stephens. 1987. K-sample Anderson-Darling tests. *Journal of the American Statistical Association* **82**(399) 918–924.
- Shaw, Aaron D, John J Horton, Daniel L Chen. 2011. Designing incentives for inexpert human raters. *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*. 275–284.

-
- Shearer, Bruce. 2004. Piece rates, fixed wages and incentives: Evidence from a field experiment. *The Review of Economic Studies* **71**(2) 513–534.
- Slakter, Malcolm J. 1965. A comparison of the Pearson chi-square and Kolmogorov goodness-of-fit tests with respect to validity. *Journal of the American Statistical Association* **60**(311) 854–858.
- Slivkins, Aleksandrs. 2019. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272* .
- Smirnov, Nickolay. 1948. Table for estimating the goodness of fit of empirical distributions. *The Annals of Mathematical Statistics* **19**(2) 279–281.
- Spear, Stephen E, Sanjay Srivastava. 1987. On repeated moral hazard with discounting. *The Review of Economic Studies* **54**(4) 599–617.
- Stephens, Michael A. 1974. Edf statistics for goodness of fit and some comparisons. *Journal of the American statistical Association* **69**(347) 730–737.
- Van der Vaart, Aad W. 2000. *Asymptotic statistics*, vol. 3. Cambridge University Press.
- Vanderbeck, François, Laurence A Wolsey. 1996. An exact algorithm for IP column generation. *Operations Research Letters* **19**(4) 151–159.
- Vera-Hernandez, Marcos. 2003. Structural estimation of a principal-agent model: Moral hazard in medical insurance. *RAND Journal of Economics* 670–693.
- Yin, Ming, Yiling Chen, Yu-An Sun. 2013. The effects of performance-contingent financial incentives in online labor markets. *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27.

Electronic Companion for “Estimation of a Non-Parametric Principal-Agent Model with Hidden Actions”

EC.1. Identifying an Optimal Incentive Contract

A natural question is whether it is possible to use our estimated model to identify an optimal incentive contract from the set R . Here, we establish an important property of our model: our agent-based method for estimating the mapping from contracts to outcomes yields an optimal contracting formulation that is simple and tractable. This tractability is a direct consequence of our specification of the agent model, and is not guaranteed if alternative methods are used to estimate the mapping from contracts to outcomes.

Let $\zeta_j(\mathbf{r})$ be the utility (e.g., of a principal) under outcome $\{\xi = j\}$ and contract \mathbf{r} , and suppose we are interested in identifying a contract $\mathbf{r} \in R$ that maximizes expected utility: $U(\mathbf{r}) = \sum_{j \in J} \Pr(\xi = j|\mathbf{r})\zeta_j(\mathbf{r})$. As we have assumed throughout, the distribution $\Pr(\xi = j|\mathbf{r})$ may be unknown in practice, which implies the utility function $U(\mathbf{r})$ is also unknown. However, given data on past contracts and outcomes, a reasonable approximation is to first estimate $\hat{\Pr}(\xi = j|\mathbf{r})$ for all $\mathbf{r} \in R$, which produces an estimate of U :

$$\hat{U}(\mathbf{r}) = \sum_{j \in J} \hat{\Pr}(\xi = j|\mathbf{r})\zeta_j(\mathbf{r}).$$

Then, a sensible approximation is to find a contract in R that maximizes the estimated expected utility:

$$(OC) \quad \underset{\mathbf{r} \in R}{\text{maximize}} \quad \hat{U}(\mathbf{r}).$$

Next, suppose $\zeta_j(\mathbf{r})$ is convex in \mathbf{r} , and let R be a convex set. Then, the tractability of the optimal contracting problem (OC) depends critically on the expression for $\hat{\Pr}(\xi = j|\mathbf{r})$, which in turn is determined by the method used to estimate it.

A useful structural property of our approach is that the expression for $\hat{\Pr}(\xi = j|\mathbf{r})$ makes problem OC quite straightforward. Specifically, within our modeling framework, we have $\hat{\Pr}(\xi = j|\mathbf{r}) = \hat{\pi}_{\hat{a}(\mathbf{r}),j}$, where $\hat{a}(\mathbf{r}) \in \operatorname{argmax}_{a \in A} \left\{ \sum_{j \in J} \hat{\pi}_{aj} r_j - c_a \right\}$. Then, for a given estimate $\hat{\boldsymbol{\pi}}$, we can write the optimal contracting problem as:

$$\underset{\mathbf{r}}{\text{maximize}} \quad \hat{\pi}_{\hat{a}(\mathbf{r}),j} \zeta_j(\mathbf{r}) \tag{EC.1a}$$

$$(OC(\hat{\boldsymbol{\pi}})) \quad \text{subject to} \quad \hat{a}(\mathbf{r}) \in \operatorname{argmax}_{a \in A} \sum_{j \in J} \hat{\pi}_{aj} r_j - c_a, \tag{EC.1b}$$

$$\mathbf{r} \in R. \tag{EC.1c}$$

This formulation leads to the following result, which we present without proof:

LEMMA EC.1. *If R is a convex set and $\zeta_j(\mathbf{r})$ is convex in \mathbf{r} , then an optimal solution to $OC(\hat{\boldsymbol{\pi}})$ can be obtained by solving $|A|$ convex problems.*

To see why the lemma holds, consider the following formulation for a fixed action a :

$$\underset{\mathbf{r}}{\text{maximize}} \quad \sum_{j \in J} \hat{\pi}_{aj} \zeta_j(\mathbf{r}) \quad (\text{EC.2a})$$

$$\text{subject to} \quad \sum_{j \in J} r_j \hat{\pi}_{aj} - c_a \geq \sum_{j \in J} r_j \hat{\pi}_{bj} - c_b, \quad b \in A, \quad (\text{EC.2b})$$

$$\mathbf{r} \in R. \quad (\text{EC.2c})$$

The subproblem (EC.2) finds a utility-maximizing contract from $R_a(\hat{\boldsymbol{\pi}})$ – the subset of R such that action a is optimal for the agent under $\hat{\boldsymbol{\pi}}$. Specifically, the objective (EC.2a) gives the expected payoff under contract \mathbf{r} and action a , the constraint (EC.2b) restricts the contracts \mathbf{r} to those that make action a optimal for the agent, and constraint (EC.2c) restricts \mathbf{r} to the set R . The subproblem (EC.2) is clearly convex if R is a convex set and $\zeta_j(\mathbf{r})$ is convex in \mathbf{r} . Therefore, $OC(\hat{\boldsymbol{\pi}})$ can be easily solved by solving the subproblem (EC.2) once for each action $a \in A$, and selecting the action and corresponding contract that maximizes (EC.2a). Note that (EC.2) may be infeasible if $R_a(\hat{\boldsymbol{\pi}})$ is empty (i.e., there exists an action such that no contract makes it optimal), but that at least one $R_a(\hat{\boldsymbol{\pi}})$ must be non-empty, so it is always possible to solve $OC(\hat{\boldsymbol{\pi}})$. In summary, once equipped with an estimate $\hat{\boldsymbol{\pi}}$, our model allows us to easily optimize over the set R to identify an optimal contract.

As a point of comparison, note that the general form optimal contracting problem OC may not be straightforward to solve if an alternative method is used to estimate $\hat{\Pr}(\xi = j|\mathbf{r})$. For example, in the context of multinomial logistic regression (where the contract vector \mathbf{r} are the independent variables and the outcomes $j \in J$ are classes), $\hat{\Pr}(\xi = j|\mathbf{r})$ is given by the logit function, which is non-convex in \mathbf{r} , making OC non-trivial to solve. Similarly, in non-parametric classification methods such as classification trees, there may not exist a closed form expression for $\hat{\Pr}(\xi = j|\mathbf{r})$, in which case solving OC is far from straightforward.

EC.2. Additional Computational Results

EC.2.1. Prediction Error Comparison with Multinomial Logistic Regression and Classification Trees

In this section, we further examine the predictive performance of the estimator using synthetic data. In particular, we assess the predictive performance of PA-D+ and the benchmark methods under different specifications of the underlying data generation process. The setup is as follows: For each observation $i \in I$, we construct \mathbf{r}^i by drawing d values from the standard uniform distribution

and sorting them in ascending order, so that $r_j^i \leq r_{j+1}^i$ for $j = 1, \dots, d-1$, which reflects the notion that higher outcomes should correspond to higher payments. For PA-D+, action costs are given by $c_a = a/m$ for $a = 1, \dots, m$. We then simulate the outcome data from discrete distributions of the form

$$\Pr(\xi = j|\mathbf{r}) = \frac{g_j(\mathbf{r})}{\sum_{k=1}^d g_k(\mathbf{r})},$$

where $g_j(\mathbf{r})$, $j \in J$ are functions that determine the outcome distribution under a contract \mathbf{r} . We test different underlying models by considering $g_j(\mathbf{r}) = r_j$, $g_j(\mathbf{r}) = \sqrt{r_j}$ and $g_j(\mathbf{r}) = (1 + r_j^2)$ (a natural interpretation is that we are testing predictive performance under different “ground-truth” models). For each of these three models, we consider six problem sizes: $(m, d, n) \in \{(2, 2, 100), (2, 2, 1000), (4, 5, 100), (4, 5, 1000), (5, 10, 100), (5, 10, 1000)\}$. For each of these six instances, we again fit all three models to a bootstrap sample of size n , and measure predictive performance on the out-of-bootstrap sample, for 100 repetitions. The prediction error is measured using the mean absolute error (MAE) given in (16). In a second set of experiments, we add a noise term ϵ to $g_j(\mathbf{r})$, where ϵ is drawn independently from the standard uniform distribution for each observation $i \in I$.

Tables 6 and 7 show the average prediction errors over 100 bootstrap repetitions with and without the additional noise term, respectively. Minimum prediction errors are shown in bold. Table 6 shows that all three methods have comparable performance when there is no noise term. In contrast, Table 7 shows that when the noise term is included to generate the outcome data, the predictive performance of MLR degrades considerably, whereas PA-D+ performs well across all instances. Intuitively, this difference in performance is due to the non-parametric nature of PA-D+, which makes predictions based on empirical distributions found within each bootstrap sample, allowing it to fit the data well. We also find that classification trees are competitive in several instances in Table 7, which is unsurprising given their flexibility.

Interestingly, PA-D+ shows little evidence of overfitting. This may be because the estimator searches over the restricted parameter set $\tilde{\Pi}$, which is constructed from the data, instead of the full space of possible distributions Π . While this restriction on the parameter space is motivated by computational tractability, it may also help the model avoid overfitting by forcing the estimator to only choose among a finite set of empirical distributions that appear in the data (see Algorithm 2). Lastly, it is worth noting that all three methods may be tuned further; for example, one can adjust the maximum number of splits in the classification trees, or adjust the parameters ρ , α and S in PA-D+; doing so may produce different results. Nonetheless, given that we do not exhaustively tune PA-D+ to the data, Tables 6 and 7 suggest that PA-D+ performs favorably compared to well-known prediction methods, and is robust to the underlying data generation process.

$g_j(\mathbf{r})$	m	d	n	PA-D+	MLR	CT
r_j	2	2	100	0.08	0.09	0.09
	2	2	1000	0.03	0.03	0.02
	4	5	100	0.07	0.08	0.07
	4	5	1000	0.02	0.03	0.02
	5	10	100	0.04	0.06	0.05
	5	10	1000	0.02	0.02	0.02
$\sqrt{r_j}$	2	2	100	0.09	0.03	0.11
	2	2	1000	0.03	0.03	0.03
	4	5	100	0.06	0.06	0.06
	4	5	1000	0.02	0.02	0.02
	5	10	100	0.04	0.05	0.05
	5	10	1000	0.02	0.02	0.01
$(1+r_j)^2$	2	2	100	0.08	0.09	0.09
	2	2	1000	0.03	0.02	0.03
	4	5	100	0.07	0.08	0.07
	4	5	1000	0.02	0.03	0.02
	5	10	100	0.05	0.05	0.06
	5	10	1000	0.02	0.02	0.02

TABLE 6. Out-of-bootstrap prediction errors of PA-D+, multinomial logistic regression (MLR) and classification trees (CT) under varying data generation processes (without noise), averaged over 100 bootstrap repetitions. Minimum errors are bolded.

$g_j(\mathbf{r})$	m	d	n	PA-D+	MLR	CT
$r_j + \epsilon$	2	2	100	0.07	0.20	0.25
	2	2	1000	0.02	0.34	0.22
	4	5	100	0.06	0.28	0.10
	4	5	1000	0.02	0.22	0.02
	5	10	100	0.04	0.16	0.05
	5	10	1000	0.02	0.13	0.02
$\sqrt{r_j} + \epsilon$	2	2	100	0.11	0.33	0.32
	2	2	1000	0.03	0.35	0.20
	4	5	100	0.07	0.27	0.08
	4	5	1000	0.02	0.22	0.02
	5	10	100	0.05	0.15	0.06
	5	10	1000	0.02	0.13	0.02
$(1+r_j)^2 + \epsilon$	2	2	100	0.10	0.19	0.31
	2	2	1000	0.03	0.23	0.20
	4	5	100	0.06	0.30	0.11
	4	5	1000	0.02	0.25	0.03
	5	10	100	0.05	0.15	0.05
	5	10	1000	0.02	0.13	0.02

TABLE 7. Out-of-bootstrap prediction errors of PA-D+, multinomial logistic regression (MLR) and classification trees (CT) under varying data generation processes (with noise), averaged over 100 bootstrap repetitions. Minimum errors are bolded.

EC.2.2. Solution Time Comparison with Maximum Likelihood Estimation

Naturally, one might ask whether π^0 can also be estimated through a maximum likelihood estimation (MLE) approach. Here, we present numerical examples to show that an MLE approach may not be tractable for the non-parametric agent model that we study in this paper, due to the potentially large search space represented by the parameter set Π and the nonconvexity of the likelihood function.

Note that for each $i \in I$, ξ^i is the outcome in J that was observed under \mathbf{r}^i . Then, based on the agent model (1), the log-likelihood of seeing the data (\mathbf{r}^i, ξ^i) , $i \in I$ under the parameter π is

$$\tilde{L}(\pi) = \sum_{i \in I} \log(\pi_{a, \xi^i}), \text{ where } a = \operatorname{argmax}_{a \in A} \left\{ \sum_{j \in J} \pi_{aj} r_j^i - c_a \right\}. \quad (\text{EC.3})$$

One challenge with maximizing the log-likelihood function $\tilde{L}(\pi)$ is that it is discontinuous at values of π where two or more actions are optimal. Intuitively, this is because for these values of π , a small change in the parameter π can make the agent's optimal action "jump" from one action to another, which leads to a discontinuity in $\tilde{L}(\pi)$. Therefore, non-linear optimization techniques that require the likelihood function to be continuous and differentiable cannot be used to maximize $\tilde{L}(\pi)$.

An intuitive approach to optimizing $\tilde{L}(\pi)$ that does not require differentiability of $\tilde{L}(\pi)$ is to perform an exhaustive grid search over the parameter set Π . We set up the following numerical experiment to test the viability of this method. First, we randomly generate problem data using the process described in §4.4.1, where $(m, d) \in \{(2, 2), (2, 3), (2, 4), (2, 5), (3, 2), (3, 3), (3, 4)\}$, and $n \in \{100, 500, 1000\}$. We then discretize the parameter set Π in increments of 0.1 for each element (a, j) in the matrix π . For example, in the instances where $d = 2$, we search over the set $\prod_{a \in A} P_{\text{grid}}$, where $P_{\text{grid}} = \{(0, 1), (0.1, 0.9), (0.2, 0.8), \dots, (0.9, 0.1), (1, 0)\}$. We then evaluate the log-likelihood function $\tilde{L}(\pi)$ for each value of π in the grid, and set the estimate as the solution that yields the largest value of $\tilde{L}(\pi)$. We set a time limit of 3600 CPU seconds for each instance.

Table 8 summarizes the results of this search procedure, averaged over 10 repetitions per problem size. The results indicate that for even small instances, an exhaustive search may require up to one CPU hour, making this method impractical for larger problem sizes. This is an unsurprising result, given that the number of solutions to be evaluated grows exponentially in the number of actions, m , and the number of outcomes, d . It can be observed from the results in Table 1 that our estimation procedure obtains comparable estimation errors in approximately 1 CPU minute for similar problem sizes. We note here that other heuristics for searching over Π (e.g. genetic algorithms) may be more fruitful than a simple grid search, although they will also be subject to a search space where the number of grid points grows exponentially in m and d .

m	d	n	MLE	
			Time	Error
2	2	100	0	0.11
2	2	500	0	0.05
2	2	1000	1	0.04
2	3	100	8	0.07
2	3	500	14	0.05
2	3	1000	22	0.04
2	4	100	144	0.07
2	4	500	264	0.05
2	4	1000	424	0.05
2	5	100	1751	0.08
2	5	500	3301	0.05
2	5	1000	-	-
3	2	100	3	0.09
3	2	500	5	0.08
3	2	1000	7	0.08
3	3	100	488	0.08
3	3	500	931	0.07
3	3	1000	1471	0.07
3	4	100	-	-
3	4	500	-	-
3	4	1000	-	-

TABLE 8. Solution time (CPU seconds) and normalized estimation error of MLE via grid search, averaged over 10 trials. Number of actions, outcomes and sample size are denoted by m , d , and n , respectively. Instances that did not solve to optimality under 3600 CPU seconds are omitted when calculating average estimation error. Dashes indicate no instance solved to optimality within 3600 CPU seconds in any trial.

EC.3. Agent Heterogeneity

The estimation procedure developed in §2 and §3 can be naturally extended to accommodate heterogeneous agents. We model heterogeneity non-parametrically by assuming each agent has a categorical *type* that is observable from the data, and allow an agent’s action costs and outcome distributions to depend on the agent’s type.

In this section, we show that our main estimator PA can be modified to incorporate agent heterogeneity, and that consistent estimation remains possible in the heterogeneous setting. In particular, we present an analogous result to Theorem 1 for an estimator that accounts for agent heterogeneity. Note that similar results to Proposition 2 and Theorem 3 can also be obtained under agent heterogeneity, by applying similar arguments. For conciseness, we will only formally present an analogue to Theorem 1 under agent heterogeneity.

We also show how *a priori* information about the relative efficiency of different agent types can be incorporated into our estimator. Note that because agent types are observable, an intuitive procedure is to simply segment the data according to agent type, and then apply the base estimator PA-D+ to each segment separately. However, if additional information about the relative performance of different agent types is available, then one might expect an estimator that *pools* the data across all agent types to outperform a naive application of PA-D+ to each type separately. We present numerical results that show that incorporating this additional efficiency information can indeed improve estimation accuracy when data is limited, but can degrade accuracy when data is abundant, due to our approximation scheme.

The extension of our model to heterogeneous agents is intuitive; therefore, the development in this section will closely follow §2 and §3. In §EC.3.1, we present the estimator formulation in the presence of agent heterogeneity, and discuss how *a priori* information about the relative efficiency of different agent types can be embedded in the parameter set Π . In §EC.3.2, we present a corollary to Theorem 1 that shows consistent estimation is possible under agent heterogeneity. In §EC.3.3, we present the optimization model and column generation algorithm used to produce an analogue to the base estimator PA-D+, and provide numerical examples. We present the proof of our main result in §B.4.

EC.3.1. Estimator

Suppose that in addition to observing the contract \mathbf{r}^i and outcome ξ^i , we also observe the agent's type, θ^i . Agent types are categorical and indexed by the set K . Thus, each historical observation consists of the triple $(\mathbf{r}^i, \xi^i, \theta^i)$. With a slight abuse of notation, we define the agent model using the parameter $\boldsymbol{\pi} \in \mathbb{R}_+^{m \times d \times |K|}$, where π_{aj}^k is the probability that action a leads to outcome j when taken by a type k agent. As before, we let $\boldsymbol{\pi}^0$ denote the true model parameter to be estimated, where $\Pi \subseteq \mathbb{R}_+^{m \times d \times |K|}$ is the parameter set and $\boldsymbol{\pi}^0 \in \Pi$. Next, let c_a^k be the cost of taking action a for a type k agent, and let \mathbf{c}^k denote a type k agent's cost vector. We assume all agents have the same outcome set, A .⁸ Let $A_k(\mathbf{r}^i, \boldsymbol{\pi}) = \operatorname{argmax}_{a \in A} \{ \sum_{j \in J} \pi_{aj}^k r_j^i - c_a^k \}$ denote the set of optimal actions for a type k agent under the contract \mathbf{r}^i , and define $I^k = \{i \in I \mid \theta^i = k\}$ as the subset of observations where the agent is type k . The loss function $L_n^k(\boldsymbol{\pi})$ for a type k agent is then

$$L_n^k(\boldsymbol{\pi}) = \underset{x, \boldsymbol{\eta}, \boldsymbol{\omega}}{\text{minimize}} \quad \|\boldsymbol{\pi} - \boldsymbol{\omega}\|_1 \tag{EC.4a}$$

$$\text{subject to } x^i \in A_k(\mathbf{r}^i, \boldsymbol{\pi}), \quad i \in I^k, \tag{EC.4b}$$

$$\omega_{aj} = \frac{1}{|\{i \in I^k \mid x^i = a\}|} \sum_{i \in \{i \in I^k \mid x^i = a\}} y_j^i, \quad a \in A, j \in J. \tag{EC.4c}$$

⁸ The assumption that all agent types share the same set of actions is without loss of generality; an action a for agent k , we may assume $c_a^k = \infty$.

Note that the loss function (EC.4) is similar to the loss function provided for a single agent type in §2.2, but is defined over the subset of observations I^k instead of I . The estimate of π^0 is then attained at a minimizer of the sum of loss functions over Π :

$$(PA-H) \quad \hat{\pi}_n = \operatorname{argmin}_{\pi \in \Pi} \sum_{k \in K} L_n^k(\pi).$$

Note that PA-H is not necessarily equivalent to solving PA for each agent type separately, because the parameter set Π may link estimation problem across agent types based on *a priori* information. In particular, if we have additional information about the relative performance of different agent types, then pooling the data and solving PA-H effectively allows us to use data generated by one agent type to estimate the parameters for another. Next, we provide two examples of how *a priori* information can be incorporated into the model.

EC.3.1.1. Outcome efficiency. It is common in incentive problems for agents to be heterogeneous with respect to some notion of *efficiency* (Laffont and Martimort 2009). Here we illustrate how different definitions of agent efficiency can be captured within our framework. Without loss of generality, assume outcomes are ordered based on the preferences of the principal, so that outcome $j + 1$ is preferred to outcome j , for $j = 1, \dots, d - 1$. Suppose also that all agents share the same action costs ($\mathbf{c}^1 = \mathbf{c}^2 = \dots, \mathbf{c}^{|K|}$). Now suppose that the outcome distributions π_a^k are unknown, except for the following (strict) first-order stochastic dominance relation: For any action a and outcome j , a type $k + 1$ agent is more likely to obtain an outcome at least as good as j when taking action a as a type k agent. Here, agent $k + 1$ can be said to be more efficient than agent k , in the sense that agent $k + 1$ is more productive under each possible action. This ordering between agent types can be captured by the following parameter set:

$$\Pi = \left\{ \pi \in Q_\pi \mid \pi \geq \mathbf{0}, \sum_{j \in J} \pi_{aj}^k = 1 \text{ for } a \in A, k \in K \right\}, \quad (EC.5)$$

where

$$Q_\pi = \left\{ \pi \mid \sum_{h=j}^d \pi_{ah}^k + \epsilon \leq \sum_{h=j}^d \pi_{ah}^{k+1}, \text{ for } a \in \{1, 2, \dots, m-1\}, j \in J, k \in \{1, 2, \dots, |K| - 1\} \right\}, \quad (EC.6)$$

where $\epsilon > 0$ is a small constant. The inequalities given in (EC.6) restrict the parameter set so that the outcome distribution is stochastically dominated when taken by a higher type (i.e., more efficient) agent. For an example where this kind of ordering may arise, consider an employee bonus program where the action set A corresponds to employee effort levels, and employees are one of two types: inexperienced ($k = 1$) or experienced ($k = 2$). Here, the outcome distribution π_a^k may be interpreted as a type k agent's productivity under action a . If the productivity of both agent types

is unknown, but there is other evidence to suggest that experienced agents are more productive than inexperienced agents (all else equal), then we might use a parameter set like the one given in (EC.5) to simultaneously estimate $\boldsymbol{\pi}^1$ and $\boldsymbol{\pi}^2$, which can be interpreted as the agents' production functions.

EC.3.1.2. Cost efficiency. Agents may also be heterogeneous with respect to action costs. Naturally, we can model heterogeneity in agent costs by assuming all agents have the same outcome distributions ($\boldsymbol{\pi}^1 = \boldsymbol{\pi}^2 = \dots = \boldsymbol{\pi}^{|K|}$), and assuming that action costs decrease in agent type: $c_a^1 \geq c_a^2 \geq \dots \geq c_a^{|K|}$, for all $a \in A$. A more general and richer model for heterogeneity in action costs is the following:

$$Q_\pi = \left\{ \boldsymbol{\pi} \left| \sum_{h=j}^d \pi_{ah}^{k+1} \leq \sum_{h=j}^d \pi_{bh}^k \text{ only if } c_b^{k+1} \geq c_a^k \right. \right\}. \quad (\text{EC.7})$$

The parameter set in (EC.7) implies that high-type agents are more cost efficient than low-type agents in the following sense: an action b taken by a low-type agent can only (weakly) stochastically dominate (i.e. be more productive than) an action a of a high-type agent if it is more costly for the low-type agent. In other words, for the low-type agent, there exists no action that is both less costly and more productive than any action of a high-type agent.

EC.3.1.3. Limiting model complexity. One potential challenge associated with our approach to incorporating agent heterogeneity is that it increases the number of parameters to be estimated; for instance, in the two examples above, the number of free parameters in the model is on the order of $m \cdot d \cdot |K|$. As with any model fitting problems, if the amount of available data is limited relative to the number of free parameters, there is a risk of the model overfitting the data, which leads to poor out-of-sample performance. A typical approach to preventing overfitting is to limit or penalize model complexity (i.e., through regularization). Within our framework, model complexity can be limited by imposing additional constraints on the parameter set Π . In particular, consider the following parameter set:

$$\Pi_{\text{reg}} = \left\{ \boldsymbol{\pi} \in \Pi \left| \sum_{(a,b) \in A \times A} \sum_{(k,k') \in K \times K} \mathbf{1}\{\boldsymbol{\pi}_a^k \neq \boldsymbol{\pi}_b^{k'}\} \leq \ell \right. \right\}, \quad (\text{EC.8})$$

where ℓ is an integer and $\ell \leq m \cdot |K|$. The set Π_{reg} may be interpreted as a regularized counterpart to Π . In words, Π_{reg} permits at most ℓ unique distributions to be used by the estimator PA-H in the construction of the matrices $\boldsymbol{\pi}^1, \boldsymbol{\pi}^2, \dots, \boldsymbol{\pi}^{|K|}$. Note that this regularization constraint can easily be imposed in the optimization formulation for PA-H (given in EC.3.3 below) through constraints on the binary assignment variables \mathbf{w} .

EC.3.2. Statistical consistency

Next, we show that the consistent estimation is possible under agent heterogeneity as well. Similar to (5), it will be helpful to define

$$R_a^k(\boldsymbol{\pi}) = \left\{ \mathbf{r} \in R \mid a \in \operatorname{argmax}_{a \in A} \sum_{j \in J} \pi_{aj}^k r_j - c_a^k \right\}.$$

Next, we present two assumptions that parallel Assumptions 1 and 2.

ASSUMPTION EC.1 (Data). *The data $(\mathbf{r}^i, \xi^i, \theta^i)$ are independent samples of random variables $(\mathbf{r}, \xi, \theta)$, where (i) $(\mathbf{r}, \xi, \theta)$ is jointly distributed with support $R \times J \times K$, (ii) \mathbf{r} has continuous marginal density function $f(\mathbf{r})$, (iii) ξ has conditional mass function $\pi_{aj}^{0k} = \Pr(\xi = j \mid \mathbf{r} \in R_a^k(\boldsymbol{\pi}))$, and (iv) for any $\boldsymbol{\pi} \in \Pi$, $\Pr(\mathbf{r} \in R_a^k(\boldsymbol{\pi})) > 0$ for all $a \in A$, $j \in J$, and $k \in K$.*

ASSUMPTION EC.2 (Identifiability). *For every $\boldsymbol{\pi} \in \Pi$ such that $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$, there exists an (a, j, k) such that*

$$\pi_{aj}^k \neq \sum_{b \in A} \pi_{bj}^{0k} \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^{0k}) \mid \mathbf{r} \in R_a(\boldsymbol{\pi}^k)).$$

We can now present a corollary to Theorem 1, which establishes that PA-H provides consistent estimates.

COROLLARY EC.1. *Let Assumption EC.1 hold. Then $\hat{\boldsymbol{\pi}}_n \rightarrow \boldsymbol{\pi}_0$ if and only if Assumption EC.2 holds.*

Corollary EC.1 confirms that the estimates produced by PA-H converge to the true model parameters. Next, we show how the integer programming formulation PA-D and the statistical column generation algorithm given in §4.3 can be extended to incorporate agent heterogeneity.

EC.3.3. Optimization, solution algorithm, and numerical examples

Similar to the single type estimator PA, we can formulate a proxy estimator for PA-H, represent this proxy estimator exactly as a mixed-integer formulation, and then obtain an integer programming formulation of a restricted estimator that limits each distribution $\boldsymbol{\pi}_a^k$ to a set of candidate distributions V . Because this development follows in a parallel manner to §3, we skip directly to the formulation of the restricted estimator.

Let \mathbf{w} , \mathbf{x} and ϕ be binary variables with the following interpretations, parallel to §4.1: $w_{as}^k = 1$ if the candidate distribution \mathbf{v}_s is assigned to distribution $\boldsymbol{\pi}_a^k$, $x_s^i = 1$ if the action that candidate distribution \mathbf{v}_s is assigned to is optimal under contract \mathbf{r}^i , and $\phi_{as}^i = 1$ if \mathbf{v}_s is assigned to distribution $\boldsymbol{\pi}_a^k$ and action a is optimal under \mathbf{r}^i and $\boldsymbol{\pi}^k$. Let $\mathbf{z}^k \in \mathbb{R}_+^{d \times S}$ be auxiliary variables. Then, analogous to PA-D, we can formulate the restricted estimator as the following optimization problem:

$$\underset{\pi, \mathbf{w}, \mathbf{x}, \mathbf{z}, \phi}{\text{minimize}} \quad \sum_{k \in K} \sum_{s \in S} \sum_{j \in J} z_{sj}^k \quad (\text{EC.9a})$$

$$\text{subject to} \quad z_{sj}^k \geq \frac{1}{n} \sum_{i \in I^k} (\xi_j^i - v_{sj}) x_s^i, \quad j \in J, s \in S, k \in K, \quad (\text{EC.9b})$$

$$z_{sj}^k \geq \frac{1}{n} \sum_{i \in I^k} (v_{sj} - \xi_j^i) x_s^i, \quad j \in J, s \in S, k \in K, \quad (\text{EC.9c})$$

$$\sum_{b \in A} \sum_{s \in S} \left(\sum_{j \in J} v_{sj} r_j^i - c_b^k \right) \phi_{bs}^i \geq \left(\sum_{j \in J} v_{s'j} r_j^i - c_a^k \right) w_{as'}^k, \quad i \in I^k, a \in A, s' \in S, k \in K, \quad (\text{EC.9d})$$

$$\text{(PA-DH)} \quad \sum_{s \in S} w_{as}^k = 1, \quad a \in A, k \in K, \quad (\text{EC.9e})$$

$$\sum_{a \in A} \sum_{s \in S} \phi_{as}^i = 1, \quad i \in I, \quad (\text{EC.9f})$$

$$x_s^i = \sum_{a \in A} \phi_{as}^i, \quad i \in I, s \in S, \quad (\text{EC.9g})$$

$$\phi_{as}^i \leq w_{as}^k, \quad i \in I^k, a \in A, s \in S, k \in K, \quad (\text{EC.9h})$$

$$\pi_{aj}^k = \sum_{s \in S} w_{as}^k v_{sj}, \quad a \in A, j \in J, k \in K, \quad (\text{EC.9i})$$

$$x_s^i \in \{0, 1\}, \quad i \in I, s \in S, \quad (\text{EC.9j})$$

$$w_{as}^k \in \{0, 1\}, \quad a \in A, s \in S, k \in K, \quad (\text{EC.9k})$$

$$\phi_{as}^i \in \{0, 1\}, \quad i \in I, a \in A, s \in S, \quad (\text{EC.9l})$$

$$\pi \in Q_\pi. \quad (\text{EC.9m})$$

The candidate distributions in the formulation above can be constructed in an analogous way to Algorithm 1. The main difference is that we first segment the data according to agent type, which is the key step in obtaining “good” candidate distributions. Recall that the intuition behind the sample-based construction of candidate distributions (given in Algorithm 1) is to use the data to approximate the true outcome distributions, π_a^0 , $a \in A$. Because each cluster in Algorithm 1 is constructed to contain *similar* contracts, we anticipate that for some clusters, all contracts that are contained within it induce the same (hidden) action from the agent. If, for a given cluster, all contracts lead to the same hidden action, then the empirical mass function over outcomes (i.e. the candidate distribution) should be informative about the outcome distribution π_a^0 for some unknown $a \in A$. In the case where agents are heterogeneous, segmenting the data by agent types before constructing the candidate distributions helps to preserve this information. A summary is given in Algorithm 3.

Next, we discuss how to use the statistical column generation algorithm to return a sufficiently “representative” subset of the candidate distributions in the presence of agent heterogeneity. Recall

Algorithm 3: Construction of candidate distributions with agent heterogeneity

Input: Data $(\mathbf{r}^i, \xi^i, \theta^i), i \in I$, parameter, $\rho > 0$.

1. For each $k \in K$, randomly sample a subset S^k from I^k .

2. **for** $s \in S^k, k \in K$:

$$B_s = \{\mathbf{r} \in R \mid \|\mathbf{r}^s - \mathbf{r}\|_2 \leq \rho\},$$

$$I_s = \{i \in I^k \mid \mathbf{r}^i \in B_s\}.$$

for $j \in J$:

$$v_{sj} = \frac{1}{|I_s|} \sum_{i \in I_s} y_j^i.$$

Output: Candidate distributions $V = \{\mathbf{v}_s \text{ for } s \in S^k, k \in K\}$.

that the statistical column generation algorithm terminates when there is no candidate distribution in the set of omitted distributions V^- that is statistically different from every distribution in V^+ . We use the same argument to build the corresponding algorithm in the presence of agent heterogeneity. We first define a vector $\boldsymbol{\psi}_s^k \in Z_+^d$, where the j^{th} entry is the frequency of outcome j in the candidate distribution \mathbf{v}_s for $s \in S^k$, obtained in Algorithm 3. Then using the test functions introduced in §4.3.1, we identify a subset of candidate distributions V^+ and solve PA-HD over V^+ instead of V . Algorithm 4 provides an overview. Note that when there is no additional structure that links the agent types, then Algorithm 3 and 4 is equivalent to running Algorithm 1 and 2 for each agent type separately.

Algorithm 4: Statistical column generation (PA-DH+)

Input: Data $(\mathbf{r}^i, \xi^i, \theta^i), i \in I$, candidate distributions V produced by Algorithm 3, significance level $\alpha > 0$.

Initialize: Set $t = 0$. Select any $s \in S^k, k \in K$. Set $S^+ = \{(s, k)\}$ and $S^- = S \setminus \{(s, k)\}$.

1. Let $(s^*, k^*) = \operatorname{argmax}_{(s, k) \in S^-} \inf_{(s', k') \in S^+} \{H_\alpha(\boldsymbol{\psi}_s^k, \boldsymbol{\psi}_{s'}^{k'})\}$.

if $\inf_{(s', k') \in S^+} \{H_\alpha(\boldsymbol{\psi}_{s^*}^{k^*}, \boldsymbol{\psi}_{s'}^{k'})\} \leq 0$ or $S^- = \emptyset$,

Solve PA-HD(S^+) and obtain solution $\boldsymbol{\pi}_n^+$, set $T = t$, and **terminate**.

else Update $t \leftarrow t + 1$, $S^+ \leftarrow \{S^+, (s^*, k^*)\}$, and $S^- \leftarrow S^- \setminus \{(s^*, k^*)\}$. Return to Step 1.

Output: Parameter estimate $\boldsymbol{\pi}_n^+$, iteration count T .

We conclude this section by providing a simple numerical example to investigate how incorporating *a priori* information about the relative efficiency of agents affects estimation error, compared to naively applying our base estimator to each agent type separately. We consider two agent types, L and H , which have true parameters $\boldsymbol{\pi}^{0L}$ and $\boldsymbol{\pi}^{0H}$, respectively. Suppose a type H agent is known to be more outcome-efficient than a type L agent (§EC.3.1.1). Then the unknown parameters $\boldsymbol{\pi}^{0L}$ and $\boldsymbol{\pi}^{0H}$ must satisfy the following inequalities, where $\epsilon > 0$ is a small constant:

$$\sum_{h=j}^d \pi_{ah}^{0L} + \epsilon \leq \sum_{h=j}^d \pi_{ah}^{0H}, \text{ for } a \in \{1, 2, \dots, m-1\}, j \in J. \quad (\text{EC.10})$$

m	d	n	Without outcome efficiency			With outcome efficiency		
			Type L	Type H	Mean	Type L	Type H	Mean
2	2	100	0.135	0.112	0.123	0.104	0.080	0.092
2	2	500	0.094	0.095	0.095	0.069	0.091	0.080
2	2	1000	0.077	0.060	0.069	0.084	0.085	0.084
4	5	100	0.058	0.052	0.055	0.054	0.043	0.049
4	5	500	0.054	0.053	0.054	0.043	0.040	0.041
4	5	1000	0.042	0.044	0.043	0.052	0.041	0.046

TABLE 9. Normalized estimation error with and without outcome efficiency constraints, averaged over 10 trials.

Recall that m and d denote the number of actions and outcomes, respectively. We consider two problem sizes, given by $(m, d) \in \{(2, 2), (4, 5)\}$. For each of the two problem sizes, we consider three sample sizes, given by $n \in \{100, 500, 1000\}$. Then for each combination (m, d, n) , we randomly generate π^{0L} and π^{0H} from Π given by (EC.5), where Q_π is given by (EC.6). We use the same data generation procedure described in §4.4.1.

Table 9 summarizes the estimation errors over 10 trials for both agent types, with and without including (EC.10) in the parameter set Q_π during estimation. Observe that the mean estimation error tends to decrease when the outcome efficiency information represented by the inequalities (EC.10) is included, when $n = 100$ and $n = 500$. Conversely, for $n = 1000$, the estimation error *increases* when the outcome efficiency information is incorporated. The intuition for this result is as follows. Note that including the constraint (EC.10) in Q_π can be interpreted as restricting the search space for the estimator LPA-DH. When n is small, the candidate distributions are poor approximations of the rows of π^{0L} and π^{0H} , and this restriction of the search space steers the estimator LPA-DH toward candidate distributions that approximate π^{0L} and π^{0H} well. However, when the sample size is large, the candidate distributions are good approximations of the rows of π^{0L} and π^{0H} , and the constraint (EC.10) restricts the search space for the estimator without providing additional information, which has the net effect of causing errors to increase.

EC.4. A Dynamic Principal-Agent Model with Hidden Actions

Our focus in this paper has been on estimating the agent model in an offline setting, where all data on historical contracts and outcomes is available at the outset, and the contract data \mathbf{r}^i , $i \in I$ is given exogenously. In this section, we consider an online counterpart to our model, where the principal can select the contracts to be offered to the agent in a dynamic manner. We note here that while extending our principal-agent framework to a dynamic setting raises new and important theoretical questions, a complete treatment of the dynamic setting is beyond the scope of this paper.

Therefore, our focus will be to present results that parallel the offline setting, namely, presenting a complete solution algorithm for the dynamic contracting problem, and proving consistency of the corresponding estimates.

The remainder of this section is organized as follows. In §C.1, we briefly discuss related literature. In §C.2, we formulate a dynamic variant of our principal-agent model. In §C.3, we present an “ ϵ -greedy” algorithm for the dynamic contracting problem, which involves iteratively solving an integer program and a sequence of linear programs. In §C.4, we present two consistency results related to the ϵ -greedy algorithm, and illustrate its performance with a simple numerical example. Proofs are contained in §C.5.

EC.4.1. Related literature

There is an extensive literature on dynamic principal-agent models with hidden actions; for examples of foundational work, see Radner (1981), Rogerson (1985), Spear and Srivastava (1987), and Abreu et al. (1990). This line of research has typically focused on characterizing the principal’s optimal decisions in an environment where parameters of the agent model are known, including the stochastic dependence of outcomes on agent actions.

The dynamic setting we consider is closer to the *multi-armed bandit problem*, which is a broad modeling framework for dynamic decision-making problems (see Slivkins (2019) for a recent review of bandit algorithms). A general setup for the multi-armed bandit problem with stochastic rewards is as follows. In each of T rounds, a utility-maximizing decision maker chooses from a set of actions (“arms”). The chosen arm generates a reward for the decision-maker, where the reward is an i.i.d. sample from an unknown distribution that depends on the arm. Because the reward distribution for each arm is unknown, the decision-maker faces what is often called an *exploration-exploitation* trade-off: they must balance exploiting “good” arms (which have been observed to produce high rewards in previous rounds) with exploring the full set of arms (which will improve knowledge about each arm’s reward distribution).

Note that our paper focuses on estimating the distribution over outcomes induced by each possible agent action. For this reason, it is natural to use the multi-armed bandit framework to formulate a dynamic variant of our principal-agent model, where each arm represents a contract, and the selected contract generates a random reward for the principal (via the agent’s hidden action). Further, unlike the offline setting, here we assume that the principal has preferences over the outcome set J . This setting is similar to work by Ho et al. (2016), who also address a dynamic principal-agent problem within a multi-armed bandit framework. Similar to our paper, the authors consider a principal-agent problem where a contract is a mapping of agent payments to outcomes, and outcomes depend on an unobservable agent action. A key distinction between Ho et al. (2016) and our work is that they assume the distribution over agent types is unknown, whereas we assume the dependence of outcomes on agent actions is the unknown parameter.

EC.4.2. Model

The principal-agent interaction proceeds over T time periods. In period t , the principal selects and offers to the agent a contract $\mathbf{r}^t \in R$. As in the offline setting, the agent selects the action that maximizes their expected utility under the selected contract. As before, the agent’s action remains hidden to the principal, who instead observes an outcome $\xi^t \in J$ in each period. If the agent takes action a , then $\xi^t = j$ with probability π_{aj}^0 , where $\boldsymbol{\pi}^0$ is unknown to the principal. Upon observing the outcome ξ^t , the principal selects a new contract \mathbf{r}^{t+1} , which initiates the next period.

Note that principal’s preferences over the outcomes J are irrelevant for estimation in the offline setting, because the contract data is already fixed. However, because in the dynamic setting the principal now chooses the contracts \mathbf{r}^t , they must also balance estimation of $\boldsymbol{\pi}^0$ with maximizing their own utility over the T rounds. Let $\zeta_j(\mathbf{r})$ be the principal’s utility under outcome $\{\xi^t = j\}$, which may include the payment r_j to the agent. Recall that $a(\mathbf{r})$ is the optimal action of the agent under \mathbf{r} and $\boldsymbol{\pi}^0$; the principal’s expected payoff under a contract \mathbf{r} is then

$$U(\mathbf{r}) = \sum_{j \in J} \pi_{a(\mathbf{r}),j}^0 \zeta_j(\mathbf{r}). \quad (\text{EC.11})$$

The principal’s dynamic contracting problem is to select the sequence of contracts $\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^T$ that maximizes their expected payoff over T periods:

$$\max_{\mathbf{r}^t \in R} \sum_{t=1}^T U(\mathbf{r}^t). \quad (\text{EC.12})$$

We shall call any \mathbf{r}^* that satisfies $\mathbf{r}^* \in \operatorname{argmax}_{\mathbf{r} \in R} U(\mathbf{r})$ an *optimal* contract. Note that the optimal solution to (EC.12) is to simply let $\mathbf{r}^t = \mathbf{r}^*$ for all $t \geq 1$. However, solving (EC.12) is challenging because the principal’s utility function $U(\mathbf{r})$ – and therefore the optimal contract \mathbf{r}^* – depends on the unknown parameter $\boldsymbol{\pi}^0$. Therefore, the principal must trade-off learning $U(\mathbf{r})$ (by estimating $\boldsymbol{\pi}^0$) with maximizing $U(\mathbf{r})$ (by selecting high-utility contracts).

EC.4.3. Algorithm overview

We now present a solution algorithm for the principal’s dynamic contracting problem (EC.12). Our approach is similar to ϵ -greedy algorithms found in the multi-armed bandit literature, which are intuitively simple and have been observed to perform well empirically (Kuleshov and Precup 2014). Within an ϵ -greedy framework, the decision-maker chooses an exploratory action in period t with probability ϵ^t , and chooses a utility-maximizing action with probability $1 - \epsilon^t$. Here, “greediness” refers to naively maximizing the decision maker’s utility using the incumbent parameter estimates. In our setting, exploration means exploring the contract set R , which is necessary for learning the parameter $\boldsymbol{\pi}^0$; exploitation means selecting the contract in R that maximizes the principal’s expected single-period utility $U(\mathbf{r})$, based on the current estimate of $\boldsymbol{\pi}^0$.

Before presenting the steps of our ϵ -greedy algorithm, we first define two new optimization problems: one for each of the exploration and exploitation steps of the algorithm.

EC.4.3.1. Exploration step. Note that the decision set to be explored is the continuous set R . To make exploration of R tractable, we first randomly sample a set of contracts \mathbf{r}_s , $s \in S$ from the contract set R , which remain fixed for the entire algorithm. Here, each \mathbf{r}_s , $s \in S$ is analogous to one “arm” in the multi-armed bandit setting. Let n_s^t be the number of times \mathbf{r}_s has been selected after the t^{th} round. Then we can construct the candidate distribution associated with \mathbf{r}^s as

$$\mathbf{v}_s^t = \frac{1}{n_s^t} \sum_{i=1}^t \mathbf{1}\{\xi^i = j\}.$$

We assume that the agent’s true optimal action $a(\mathbf{r}^s)$ is unique under each \mathbf{r}^s (note that this is almost surely the case if the initial arms \mathbf{r}^s , $s \in S$ are selected by randomly sampling from R according to a continuous distribution). Each \mathbf{v}_s^t can then be interpreted as an empirical distribution constructed by sampling n_s^t times from the distribution $\pi_{a(\mathbf{r}^s)}^0$.

Next, define a set of error variables $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_{|S|}$, and let w_{as} be a binary variable equal to 1 if \mathbf{v}_s is assigned to $a \in A$. We now define the following optimization problem:

$$\underset{w, \varepsilon}{\text{minimize}} \quad \sum_{s \in S} |\varepsilon_s| \tag{EC.13a}$$

$$\text{subject to} \quad \sum_{a \in A} \left(\sum_{j \in J} v_{s'j}^t r_{sj} - c_a \right) w_{as} + \varepsilon_s \geq \sum_{a \in A} \left(\sum_{j \in J} v_{s'j}^t r_{sj} - c_a \right) w_{as'}, \quad s \in S, s' \in S, \tag{EC.13b}$$

$$\sum_{a \in A} w_{as} = 1, \quad s \in S, \tag{EC.13c}$$

$$\text{(PA-T)} \quad \sum_{s \in S} w_{as} \geq 1, \quad a \in A, \tag{EC.13d}$$

$$w_{as} \in \{0, 1\}, \quad a \in A, s \in S. \tag{EC.13e}$$

$$\varepsilon_s \geq 0, \quad s \in S. \tag{EC.13f}$$

Intuitively, formulation PA-T assigns each contract \mathbf{r}^s to the action $a \in A$ that is believed to be optimal under \mathbf{r}^s and the true model $\boldsymbol{\pi}^0$. Constraint (EC.13c) ensures every contract \mathbf{r}^s is assigned to exactly one action, and (EC.13d) ensures every action has at least one contract assigned to it. The key constraint of this formulation is (EC.13b): Given a contract \mathbf{r}^s , constraint (EC.13b) forces the action associated with \mathbf{r}^s to yield at least as large of a utility for the agent as the action associated with $\mathbf{r}^{s'}$, for every $s' \in S$. Note that because \mathbf{v}_s^t is an approximation of $\boldsymbol{\pi}_a^0$ for some $a \in A$, constraint (EC.13b) may not hold when $\varepsilon_s = 0$, even if the binary decision variable \mathbf{w} correctly assigns each \mathbf{r}^s to the agent’s true optimal action. Therefore, ε_s is required to serve as a slack variable that maintains feasibility of PA-T. Intuitively, ε_s can be interpreted as measuring the sub-optimality of the assignment encoded by \mathbf{w} .

Let $(\bar{\mathbf{w}}^t, \bar{\boldsymbol{\epsilon}}^t)$ be the optimal solution to PA-T in round t . Then we construct the estimate $\hat{\boldsymbol{\pi}}^t$ as follows:

$$\hat{\pi}_{aj}^t = \frac{\sum_{s \in S} v_{sj}^t n_s^t \bar{w}_{as}^t}{\sum_{s \in S} n_s^t \bar{w}_{as}^t}, \quad \text{for } a \in A, j \in J.$$

Intuitively, $\hat{\pi}_a^t$ is the average of all candidate distributions assigned to action a , weighted by their sample sizes n_s^t .

EC.4.3.2. Exploitation step. The exploitation step consists of solving the optimal contracting problem $\text{OC}(\hat{\boldsymbol{\pi}}^t)$, described in §EC.1 of the electronic companion, under the incumbent estimate $\hat{\boldsymbol{\pi}}^t$.

EC.4.3.3. Algorithm summary. Having defined the optimization models PA-C and OC, we can now summarize the steps of the ϵ -greedy algorithm (Algorithm 5). At each iteration, an exploration action is taken with probability ϵ^t , by randomly selecting one of the contracts from $\{\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^{|S|}\}$ and updating the set of candidate distributions accordingly, and an exploitation action is taken with probability $1 - \epsilon^t$, by selecting the optimal contract based on the current estimate, $\hat{\boldsymbol{\pi}}$.

ALGORITHM 5: ϵ -greedy algorithm for dynamic contracting

Input: Exploration parameters $\epsilon^1, \epsilon^2, \dots, \epsilon^T$.

1. **Initialize:** Randomly sample S contracts $\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^S$ from set R . Initialize $n_s^t = 1$ for $s \in S$ and $t = 1$.

2. **for** $t = 1, 2, \dots, T$:

 with probability ϵ^t :

 Set $\hat{\mathbf{r}}^t = \hat{\mathbf{r}}^{t-1}$.

 Randomly select $\bar{s} \in S$. Select contract $\mathbf{r}^{\bar{s}}$ and observe outcome $\bar{\xi}$. Set $s^t = \bar{s}$ and $\xi^t = \bar{\xi}$.

 Update $n_{\bar{s}}^t \leftarrow n_{\bar{s}}^t + 1$ and $v_{\bar{s}j}^t = \frac{1}{n_{\bar{s}}^t} \sum_{i=1}^t \mathbf{1}\{\xi^i = j, s^i = \bar{s}\}$ for $j \in J$.

 with probability $1 - \epsilon^t$:

 Solve PA-T and obtain $\bar{\mathbf{w}}^t$. For each $a \in A, j \in J$, set $\hat{\pi}_{aj}^t = \frac{\sum_{s \in S} v_{sj}^t n_s^t \bar{w}_{as}^t}{\sum_{s \in S} n_s^t \bar{w}_{as}^t}$.

 Solve $\text{OC}(\hat{\boldsymbol{\pi}}^t)$ and obtain $\hat{\mathbf{r}}^t$.

Output: Estimates $\hat{\boldsymbol{\pi}}^T$ and $\hat{\mathbf{r}}^T$.

EC.4.4. Consistency, numerical examples, and discussion

We now present the main theoretical result of this section. First, consider the following assumption:

ASSUMPTION EC.3. *For each $s \in S$, there exists $s' \in S$ such that*

$$\sum_{j \in J} (\pi_{a(\mathbf{r}^{s'}), j}^0 - \pi_{a(\mathbf{r}^s), j}^0) r_j^s > \bar{c} - \underline{c}, \quad (\text{EC.14})$$

where $\bar{c} = \sup_{a \in A} \{c_a\}$ and $\underline{c} = \inf_{a \in A} \{c_a\}$.

Assumption EC.3 is an identifiability condition that ensures precise inference of $\boldsymbol{\pi}^0$, similar to Assumption 2 in the offline setting. Loosely speaking, this condition ensures that if there exists a solution $\bar{\mathbf{w}}$ to PA-T that assigns a candidate distribution v_s to the wrong action, then the assignments encoded in $\bar{\mathbf{w}}$ will eventually (as more data is collected) violate the optimality conditions in (EC.13b). This has the effect of guaranteeing that the solution to PA-T identifies the correct action for each candidate distribution v_s in the limit. Next, Proposition EC.1 states that the condition in EC.3 is sufficient for Algorithm 5 to uncover the true values of $\boldsymbol{\pi}^0$.

PROPOSITION EC.1. *Let Assumption EC.3 hold. Then the estimate $\hat{\boldsymbol{\pi}}^T$ produced by Algorithm 5 is consistent:*

$$\hat{\boldsymbol{\pi}}^T \longrightarrow \boldsymbol{\pi}^0.$$

Note that Assumption EC.3 is stronger than the analogous condition in the offline setting, Assumption EC.2). The reason for this difference is as follows: In the online setting, we restrict attention to the subset of contracts $\{\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^{|S|}\}$ when estimating $\hat{\boldsymbol{\pi}}$. While focusing on this discrete subset of R improves tractability, it also reduces variation in the data compared to the offline setting, which reduces the information available to the estimator. Because the online setting has less variation in contract data, a stronger identifiability condition is required to ensure that learning the precise values of the parameter $\boldsymbol{\pi}^0$ remains possible.

COROLLARY EC.2. *Let Assumption EC.3 hold. Then the contract $\hat{\mathbf{r}}^T$ produced by Algorithm 5 converges to a minimizer of the principal’s per-period regret:*

$$\left| \frac{1}{T} \sum_{t=1}^T U(\hat{\mathbf{r}}^T) - \frac{1}{T} \sum_{t=1}^T U(\mathbf{r}^*) \right| \longrightarrow 0.$$

The intuition behind Corollary EC.2 is straightforward – because $\hat{\boldsymbol{\pi}}^T$ is a consistent estimate of $\boldsymbol{\pi}^0$, then solving $\text{OC}(\hat{\boldsymbol{\pi}}^T)$ will eventually produce an optimal contract for the principal, as $T \rightarrow \infty$.

We now illustrate Algorithm 5 through numerical examples. We consider two problem sizes, given by $(m, d) \in \{(2, 2), (5, 10)\}$. For each problem size, we construct the contract set as $R = [1, 2]^d$, the agent cost vector \mathbf{c} as a random sample from $[0, 1]^m$. We specify the principal’s utility as $\zeta_j(\mathbf{r}) = \bar{\zeta}_j - r_j$, where $\bar{\zeta}_j$ is a random sample from $[1, 10]^d$. We construct the true parameter $\boldsymbol{\pi}^0$ by letting $\boldsymbol{\pi}_a^0$ for each $a \in A$ be randomly generated from the $(d - 1)$ -dimension simplex. We test six parameterizations of the exploration probability ϵ^t . We first consider three “fixed” exploration schemes, where $\epsilon^t = \epsilon$, for each $\epsilon \in \{0.5, 0.9, 0.99\}$. We then consider three “variable” exploration schemes, where $\epsilon^t = \exp(-\lambda \cdot t)$, for each $\lambda \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. For each problem size and parameterization of ϵ^t , we run 10 trials of Algorithm 5 with $T = 1000$, where $\boldsymbol{\pi}^0$, $\bar{\boldsymbol{\zeta}}$, and \mathbf{c} are randomly constructed in each trial.

m	d	T	$\epsilon^t = \epsilon$			$\epsilon^t = \exp(-\lambda \cdot t)$		
			$\epsilon = 0.5$	$\epsilon = 0.9$	$\epsilon = 0.99$	$\lambda = 10^{-2}$	$\lambda = 10^{-3}$	$\lambda = 10^{-4}$
2	2	50	0.21	0.11	0.07	0.06	0.06	0.06
2	2	100	0.20	0.11	0.07	0.06	0.06	0.06
2	2	500	0.16	0.08	0.07	0.05	0.06	0.06
2	2	1000	0.13	0.08	0.07	0.06	0.06	0.06
2	2	50	0.12	0.16	0.14	0.10	0.14	0.16
2	2	100	0.09	0.16	0.15	0.09	0.14	0.16
2	2	500	0.08	0.14	0.16	0.03	0.11	0.16
2	2	1000	0.08	0.14	0.16	0.02	0.09	0.16

TABLE 10. Estimation errors (top) and per-period regret (bottom) for six exploration schemes, averaged over 10 trials ($m = 2, d = 2$).

m	d	T	$\epsilon^t = \epsilon$			$\epsilon^t = \exp(-\lambda \cdot t)$		
			$\epsilon = 0.5$	$\epsilon = 0.9$	$\epsilon = 0.99$	$\lambda = 10^{-2}$	$\lambda = 10^{-3}$	$\lambda = 10^{-4}$
5	10	50	0.05	0.06	0.06	0.05	0.05	0.05
5	10	100	0.06	0.06	0.06	0.05	0.05	0.05
5	10	500	0.06	0.06	0.05	0.05	0.05	0.05
5	10	1000	0.06	0.06	0.05	0.05	0.05	0.05
5	10	50	0.78	0.82	0.87	0.78	0.84	0.85
5	10	100	0.77	0.83	0.87	0.74	0.83	0.84
5	10	500	0.78	0.82	0.85	0.58	0.79	0.83
5	10	1000	0.79	0.82	0.85	0.56	0.73	0.83

TABLE 11. Estimation errors (top) and per-period regret (bottom) for six exploration schemes, averaged over 10 trials ($m = 5, d = 10$).

A summary of results is provided in Table 10 ($m = 2, d = 2$) and Table 11 ($m = 5, d = 10$). For each of the six exploration schemes, we report the average estimation error and average per-period regret at $t = 100$, $t = 500$ and $t = 1000$. Among the fixed exploration schemes, aggressive exploitation ($\epsilon^t = 0.5$) attains the lowest regret and highest estimation error, and aggressive exploration ($\epsilon^t = 0.99$) attain the highest regret and lowest estimation error. These results demonstrate the exploration-exploitation trade-off in our dynamic contracting problem. Comparing all six exploration schemes, the variable scheme with $\lambda = 10^{-2}$ appears to weakly dominate, by producing similar estimation errors as the other approaches, but with significantly lower regret. This suggests that $\lambda = 10^{-2}$ handles the exploration-exploitation trade-off the most efficiently. We also observe that the decrease in regret is significantly more pronounced for the variable scheme with $\lambda = 10^{-2}$ compared to the other approaches, because it exploits the most aggressively as t increases (e.g., by $t = 500$, this scheme exploits with probability 0.99.)

We conclude by highlighting directions for future work. First, the algorithm we presented in this section is based on ϵ -greedy approaches found in the multi-armed bandit literature. Other

well-known bandit algorithms may also be applicable in our setting, such as those based on upper confidence bounds or Thompson sampling (Slivkins 2019). Second, while we have focused our analysis in this section on asymptotic results, the standard performance measure of a dynamic decision-making algorithm is a finite-sample bound on regret. It may be possible to obtain similar bounds in our dynamic contracting setting. However, doing so will require addressing two technical challenges of our setting that are not present in classical bandit problems: the continuous nature of the decision space R (instead of a discrete set of arms), and discontinuity in the principal's payoff in \mathbf{r} (due to jumps in the agent's optimal action as \mathbf{r} changes). Third, we have assumed throughout that the agent behaves myopically by optimizing their single-period payoff. It may also be the case that agents behave strategically, which would introduce new and challenging dynamics into the problem.

Note also that the performance of Algorithm 5 is sensitive to how R is specified. Because R is the exploration space, if the values in R are large compared to the principal's utility, then the algorithm may spend a large number of iterations exploring regions of R that are inefficient from the principal's perspective. This sensitivity to the contract set R occurs because the Algorithm 5 does not consider regret when selecting a contract to explore, which is common in ϵ -greedy approaches.

EC.5. Proofs

Because the proof for Theorem 2 is long, we group all proofs into four subsections for ease of navigation: §EC.5.1 contains the proofs for Theorem 1 and Proposition 1; §EC.5.2 contains the proof of Theorem 2; §EC.5.3 contains the proofs of Proposition 2, Theorem 3, and Theorem 4; and §E.4 contains the proofs for Sections EC.3 and EC.4.

EC.5.1. Proofs of Theorem 1 and Proposition 1

Before proving Theorem 1, we prove two supporting results, given in Lemmas EC.2 and EC.3. Lemma EC.2 shows that the loss function $L_n(\boldsymbol{\pi})$ is lower semicontinuous. Lemma EC.3 shows that as $n \rightarrow \infty$, the loss function $L_n(\boldsymbol{\pi})$ converges point-wise to a function $L(\boldsymbol{\pi})$, where $L(\boldsymbol{\pi})$ is uniquely minimized by the true parameter $\boldsymbol{\pi}^0$ if and only if Assumption 2 holds. Both of these results are used to prove the consistency result in Theorem 1.

LEMMA EC.2. $L_n(\boldsymbol{\pi})$ is lower semicontinuous in $\boldsymbol{\pi}$ on Π for all $n \geq 1$.

Proof. Define $N(\delta) = \{\boldsymbol{\pi} \in \Pi \mid \|\boldsymbol{\pi} - \bar{\boldsymbol{\pi}}\|_1 < \delta\}$. To show that $L_n(\boldsymbol{\pi})$ is lower semicontinuous at all $\bar{\boldsymbol{\pi}} \in \Pi$ and $n \geq 1$, it suffices to prove the following statement (Rockafellar and Wets 2009): For every $\varepsilon > 0$, there exists $\delta > 0$ such that $L_n(\bar{\boldsymbol{\pi}}) - L_n(\boldsymbol{\pi}) < \varepsilon$ for all $\boldsymbol{\pi} \in N(\delta)$. The proof proceeds

in two steps. First, we show that there exists $\bar{\delta} > 0$ such that for all $\delta \in (0, \bar{\delta})$ and $\boldsymbol{\pi} \in N(\delta)$, $A(\mathbf{r}^i, \boldsymbol{\pi}) \subseteq A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$ for all $i \in I$. Second, we show lower semicontinuity of $L_n(\boldsymbol{\pi})$ at $\bar{\boldsymbol{\pi}}$. Step 1. Fix $\bar{\boldsymbol{\pi}}$ and n . By way of contradiction, suppose that for all $\delta > 0$, there exists $\check{\boldsymbol{\pi}} \in N(\delta)$, $a \in A(\mathbf{r}^i, \check{\boldsymbol{\pi}})$ and $i \in I$ such that $a \notin A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$. Note that $a \notin A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$ and $a \in A(\mathbf{r}^i, \check{\boldsymbol{\pi}})$ implies there exists a $b \in A \setminus a$ such that the following two inequalities hold:

$$\sum_{j \in J} \bar{\pi}_{aj} r_j^i - c_a < \sum_{j \in J} \bar{\pi}_{bj} r_j^i - c_b \quad (\text{EC.15})$$

$$\sum_{j \in J} \check{\pi}_{aj} r_j^i - c_a \geq \sum_{j \in J} \check{\pi}_{bj} r_j^i - c_b. \quad (\text{EC.16})$$

Combining (EC.15) and (EC.16) yields

$$\sum_{j \in J} (\check{\pi}_{aj} - \bar{\pi}_{aj}) r_j^i - c_a > \sum_{j \in J} (\check{\pi}_{bj} - \bar{\pi}_{bj}) r_j^i - c_b.$$

Letting $\delta \rightarrow 0$ implies $\|\check{\boldsymbol{\pi}} - \bar{\boldsymbol{\pi}}\|_1 \rightarrow 0$ and thus $0 > 0$, a contradiction. It follows that there exists $\bar{\delta} > 0$ such that for all $\delta \in (0, \bar{\delta})$ and $\boldsymbol{\pi} \in N(\delta)$, $A(\mathbf{r}^i, \boldsymbol{\pi}) \subseteq A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$ for all $i \in I$. Step 2. Fix $\varepsilon > 0$. By way of contradiction, suppose $L_n(\boldsymbol{\pi})$ is not lower semicontinuous; that is, for every $\delta > 0$, there exists $\check{\boldsymbol{\pi}}_\delta \in N(\delta)$ such that $L_n(\bar{\boldsymbol{\pi}}) - L_n(\check{\boldsymbol{\pi}}_\delta) > \varepsilon$. Next, let $(\mathbf{x}(\boldsymbol{\pi}), \boldsymbol{\omega}(\boldsymbol{\pi}))$ be an optimal solution to (4) under $\boldsymbol{\pi} = \check{\boldsymbol{\pi}}_\delta$. By Step 1, if $\delta \in (0, \bar{\delta})$, then $A(\mathbf{r}^i, \check{\boldsymbol{\pi}}_\delta) \subseteq A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$ for all $i \in I$. Because $A(\mathbf{r}^i, \check{\boldsymbol{\pi}}_\delta) \subseteq A(\mathbf{r}^i, \bar{\boldsymbol{\pi}})$ for all $i \in I$, it is feasible to set $\boldsymbol{\pi} = \bar{\boldsymbol{\pi}}$ and $(\mathbf{x}, \boldsymbol{\omega}) = (\mathbf{x}(\check{\boldsymbol{\pi}}_\delta), \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta))$ in (4). It follows that $L_n(\bar{\boldsymbol{\pi}}) = \|\bar{\boldsymbol{\pi}} - \boldsymbol{\omega}(\bar{\boldsymbol{\pi}})\|_1 \leq \|\bar{\boldsymbol{\pi}} - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1$. Therefore,

$$\varepsilon < L_n(\bar{\boldsymbol{\pi}}) - L_n(\check{\boldsymbol{\pi}}_\delta) \leq \|\bar{\boldsymbol{\pi}} - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1 - L_n(\check{\boldsymbol{\pi}}_\delta) = \|\bar{\boldsymbol{\pi}} - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1 - \|\check{\boldsymbol{\pi}} - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1.$$

Further, it is straightforward to verify that

$$\|\bar{\boldsymbol{\pi}} - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1 - \|\check{\boldsymbol{\pi}}_\delta - \boldsymbol{\omega}(\check{\boldsymbol{\pi}}_\delta)\|_1 \leq \|\bar{\boldsymbol{\pi}} - \check{\boldsymbol{\pi}}_\delta\|_1.$$

Thus, $\varepsilon \leq \|\bar{\boldsymbol{\pi}} - \check{\boldsymbol{\pi}}_\delta\|_1$. Letting $\delta \rightarrow 0$ implies $\|\bar{\boldsymbol{\pi}} - \check{\boldsymbol{\pi}}_\delta\|_1 \rightarrow 0$, which yields a contradiction. We conclude that for every $\varepsilon > 0$, there exists $\delta > 0$ such that $L_n(\bar{\boldsymbol{\pi}}) - L_n(\boldsymbol{\pi}) < \varepsilon$ for all $\boldsymbol{\pi} \in N(\delta)$, and thus $L_n(\boldsymbol{\pi})$ is lower semicontinuous. \square

LEMMA EC.3. *There exists $L(\boldsymbol{\pi}) : \Pi \rightarrow \mathbb{R}$ such that (i) $L(\boldsymbol{\pi}) < \infty$ and $|L_n(\boldsymbol{\pi}) - L(\boldsymbol{\pi})| \rightarrow 0$ for all $\boldsymbol{\pi} \in \Pi$, (ii) $L(\boldsymbol{\pi}^0) = 0$, and (iii) $L(\boldsymbol{\pi}^0) < L(\boldsymbol{\pi})$ for every $\boldsymbol{\pi} \in \Pi$ such that $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$ if and only if Assumption 2 holds.*

Proof. The proof proceeds in two steps. In Step 1, we show the following supporting result: For any (a, j) ,

$$\lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi}) = \frac{1}{\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}))} \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \quad (\text{EC.17})$$

where $R_a(\boldsymbol{\pi})$ is defined in (5). In Step 2, we prove there exists a function $L(\boldsymbol{\pi})$ such that statements (i), (ii) and (iii) hold. Step 1. Pick any (a, j) . Note that $x^i = a$ holds in (4) holds if and only if $\mathbf{r}^i \in R_a(\boldsymbol{\pi})$. For each (a, j) , we can then write ω_{aj} as

$$\omega_{aj}(\boldsymbol{\pi}) = n \cdot \frac{1}{n} \cdot \omega_{aj}(\boldsymbol{\pi}) = \frac{n}{\sum_{i=1}^n \mathbb{I}\{\mathbf{r}^i \in R_a(\boldsymbol{\pi})\}} \left(\frac{1}{n} \sum_{i \in \{i | \mathbf{r}^i \in R_a(\boldsymbol{\pi})\}} y_j^i \right). \quad (\text{EC.18})$$

Consider the first term on the right-hand side of (EC.18). By Assumption 1, letting $n \rightarrow \infty$ yields

$$\frac{n}{\sum_{i=1}^n \mathbb{I}\{\mathbf{r}^i \in R_a(\boldsymbol{\pi})\}} = \frac{1}{\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}))}. \quad (\text{EC.19})$$

For the second term on the right-hand side of (EC.18), we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in \{i | \mathbf{r}^i \in R_a(\boldsymbol{\pi})\}} y_j^i &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{\xi^i = j, \mathbf{r}^i \in R_a(\boldsymbol{\pi})\}, \\ &= \Pr(\xi = j, \mathbf{r} \in R_a(\boldsymbol{\pi})), \\ &= \sum_{b \in A} \Pr(\xi = j, \mathbf{r} \in R_a(\boldsymbol{\pi}) | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)) \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^0)). \end{aligned}$$

The first equality above follows because $y_j^i = \mathbb{I}\{\xi^i = j\}$, by definition. The second equality follows from the strong law of large numbers. The third equality follows from the law of total probability. Next, note that by Assumption 1, the events $\{\xi = j\}$ and $\{\mathbf{r} \in R_a(\boldsymbol{\pi})\}$ are conditionally independent given $\{\mathbf{r} \in R_b(\boldsymbol{\pi}^0)\}$. Therefore,

$$\begin{aligned} \Pr(\xi = j, \mathbf{r} \in R_a(\boldsymbol{\pi}) | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)) &= \Pr(\xi = j | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)) \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}) | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \\ &= \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}) | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \end{aligned}$$

where the second equality follows by definition of $\boldsymbol{\pi}^0$. It follows that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in \{i | \mathbf{r}^i \in R_a(\boldsymbol{\pi})\}} y_j^i &= \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}) | \mathbf{r} \in R_b(\boldsymbol{\pi}^0)) \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \\ &= \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)). \end{aligned} \quad (\text{EC.20})$$

Next, combining (EC.18), (EC.19), and (EC.20) yields

$$\lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi}) = \frac{1}{\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}))} \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)),$$

as desired. Step 2. Define $L(\boldsymbol{\pi}) = \sum_{a \in A} \sum_{j \in J} |\pi_{aj} - \lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi})|$. We prove (i), (ii) and (iii) in order. (i). Note by continuity of the absolute value function, $\lim_{n \rightarrow \infty} |\pi_{aj} - \omega_{aj}(\boldsymbol{\pi})| = |\pi_{aj} - \lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi})|$ for all (a, j) and $\boldsymbol{\pi} \in \Pi$. It follows by definition of $L(\boldsymbol{\pi})$ that $|L_n(\boldsymbol{\pi}) - L(\boldsymbol{\pi})| \rightarrow 0$. Further, $L(\boldsymbol{\pi}^0) < \infty$ follows by definition of $L(\boldsymbol{\pi})$ and because $\lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi}) < \infty$ by Step 1. (ii).

By definition of $L(\boldsymbol{\pi})$, to show that $L(\boldsymbol{\pi}^0) = 0$, it suffices to show that $\pi_{aj}^0 = \lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi}^0)$ for all (a, j) . Pick any (a, j) , and note

$$\begin{aligned} \pi_{aj}^0 &= \frac{1}{\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0))} \cdot \pi_{aj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0)), \\ &= \frac{1}{\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0))} \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \\ &= \lim_{n \rightarrow \infty} \omega_{aj}(\boldsymbol{\pi}^0), \end{aligned}$$

where the first equality follows by multiplying and dividing by $\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0))$, the second equality follows because $\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi}^0), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)) = 0$ for all $b \neq a$ by continuity of $f(\mathbf{r})$, and the third equality follows from Step 1. Therefore, $L(\boldsymbol{\pi}^0) = 0$. (iii). Next, we establish that $L(\boldsymbol{\pi}) > 0$ for all $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$ if and only if Assumption 2 holds. First, let Assumption 2 hold. Pick any $\bar{\boldsymbol{\pi}} \neq \boldsymbol{\pi}^0$, and suppose by way of contradiction that $L(\bar{\boldsymbol{\pi}}) = 0$. It follows that for all (a, j) ,

$$\begin{aligned} \bar{\pi}_{aj} &= \lim_{n \rightarrow \infty} \omega_{aj}(\bar{\boldsymbol{\pi}}), \\ &= \frac{1}{\Pr(\mathbf{r} \in R_a(\bar{\boldsymbol{\pi}}))} \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_a(\bar{\boldsymbol{\pi}}), \mathbf{r} \in R_b(\boldsymbol{\pi}^0)), \\ &= \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^0) | \mathbf{r} \in R_a(\bar{\boldsymbol{\pi}})), \end{aligned}$$

where the first line follows because $L(\bar{\boldsymbol{\pi}}) = 0$, the second line follows from Step 1, and the third line follows by the probability chain rule. However, by Assumption 2, there exists an (a, j) such that $\bar{\pi}_{aj} \neq \sum_{b \in A} \pi_{bj}^0 \cdot \Pr(\mathbf{r} \in R_b(\boldsymbol{\pi}^0) | \mathbf{r} \in R_a(\bar{\boldsymbol{\pi}}))$, which yields a contradiction. Therefore, $L(\boldsymbol{\pi}) > 0$ for all $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$. Conversely, if Assumption 2 does not hold, then by parallel argument to the above, there exists $\check{\boldsymbol{\pi}} \in \Pi$ where $\check{\boldsymbol{\pi}} \neq \boldsymbol{\pi}^0$ such that $\check{\pi}_{aj} = \lim_{n \rightarrow \infty} \omega_{aj}(\check{\boldsymbol{\pi}})$ for all (a, j) . By definition of $L(\boldsymbol{\pi})$, it follows that $L(\boldsymbol{\pi}^0) = L(\check{\boldsymbol{\pi}}) = 0$. \square

Proof of Theorem 1. The proof proceeds in two steps. First, we show $|L_n(\hat{\boldsymbol{\pi}}_n) - L_n(\boldsymbol{\pi}^0)| \rightarrow 0$. Second, we show $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n = \boldsymbol{\pi}^0$ if and only if Assumption 2 holds. Step 1. Because $L_n(\boldsymbol{\pi}^0) \rightarrow L(\boldsymbol{\pi}^0)$ by Lemma EC.3(i) and $L(\boldsymbol{\pi}^0) = 0$ by Lemma EC.3(ii), we have $L_n(\boldsymbol{\pi}^0) \rightarrow 0$. Next, note

$$0 \leq L_n(\hat{\boldsymbol{\pi}}_n) \leq L_n(\boldsymbol{\pi}^0),$$

where the first and second inequalities follow by definition of $L_n(\boldsymbol{\pi})$ and $\hat{\boldsymbol{\pi}}_n$, respectively. It follows that $L_n(\hat{\boldsymbol{\pi}}_n) \rightarrow 0$. Therefore, $|L_n(\hat{\boldsymbol{\pi}}_n) - L_n(\boldsymbol{\pi}^0)| \rightarrow 0$. Step 2. Because $L_n(\boldsymbol{\pi})$ is lower semicontinuous by Lemma EC.2, $|L_n(\hat{\boldsymbol{\pi}}_n) - L_n(\boldsymbol{\pi}^0)| \rightarrow 0$ by Step 1, and Π is compact, by Theorem 5.14 of Van der Vaart (2000), $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n \in \text{argmin}_{\boldsymbol{\pi} \in \Pi} L(\boldsymbol{\pi})$. Next, suppose Assumption 2 holds. Then by Lemma EC.3, $\text{argmin}_{\boldsymbol{\pi} \in \Pi} L(\boldsymbol{\pi}) = \boldsymbol{\pi}^0$, which implies $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n = \boldsymbol{\pi}^0$. If Assumption 2

does not hold, then by Lemma EC.3(iii), there exists $\tilde{\boldsymbol{\pi}} \neq \boldsymbol{\pi}^0$ such that $L(\tilde{\boldsymbol{\pi}}) = L(\boldsymbol{\pi}^0) = 0$, in which case $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n \neq \boldsymbol{\pi}^0$. \square

Proof of Proposition 1. We prove statements (i), (ii) and (iii) in order. Part (i). To show that the minimizer of PA-C also minimizes the proxy loss function $Z_n(\boldsymbol{\pi})$, it suffices to show that solving PA-C and minimizing $Z_n(\boldsymbol{\pi})$ are equivalent problems in the following sense: (a) For any $(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ that is feasible for PA-C, $\boldsymbol{\pi} \in \Pi$; (b) For any $\boldsymbol{\pi} \in \Pi$, there exists (\mathbf{x}, \mathbf{z}) such that $(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ is feasible to PA-C; and (c) for any $(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ feasible for PA-C, $Z_n(\boldsymbol{\pi}) = Z_n^C(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$, where $Z_n^C(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ is the objective of PA-C. Statements (a) and (b) follow immediately by the construction of PA-C and the proxy loss problem (8); it remains to prove (c). Let $\boldsymbol{\pi}$ be fixed in PA-C, and let (\mathbf{x}, \mathbf{z}) be a solution to the resulting subproblem. Next, note

$$\begin{aligned} \sum_{i \in I} (\pi_{aj} - y_j^i) x_a^i &= \sum_{i \in \{i | x^i = a\}} (\pi_{aj} - y_j^i), \\ &= (\pi_{aj} - \omega_{aj}) |\{i | x^i = a\}|, \\ &= (\pi_{aj} - \omega_{aj}) \eta_{aj}, \end{aligned} \tag{EC.21}$$

where the first equality follows because $x_a^i = 1$ if and only if $i \in \{i | x^i = a\}$, the second equality follows from the definition of $\boldsymbol{\omega}$, and the third equality follows by definition of $\boldsymbol{\eta}$. Then we have

$$\begin{aligned} Z_n^C(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z}) &= \sum_{a \in A} \sum_{j \in J} \left| \frac{1}{n} \sum_{i \in I} (\pi_{aj} - y_j^i) x_a^i \right|, \\ &= \frac{1}{n} \sum_{a \in A} \sum_{j \in J} |(\pi_{aj} - \omega_{aj}) \eta_{aj}|, \\ &= Z_n(\boldsymbol{\pi}), \end{aligned}$$

where the first equality follows from (10a)–(10c), the second equality follows from (EC.21), the third equality follows by definition of the element-wise norm $\|\cdot\|_1$, and the final equality follows by definition of $Z_n(\boldsymbol{\pi})$. Therefore, $Z_n(\boldsymbol{\pi}) = Z_n^C(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ for any $(\boldsymbol{\pi}, \mathbf{x}, \mathbf{z})$ feasible for PA-C. It follows that $\boldsymbol{\pi}_n^* \in \text{argmin}_{\boldsymbol{\pi} \in \Pi} Z_n(\boldsymbol{\pi})$. Part (ii). The proof of Theorem 1 establishes that $L_n(\hat{\boldsymbol{\pi}}_n) \rightarrow 0$. Note $L_n(\hat{\boldsymbol{\pi}}_n) \leq L_n(\bar{\boldsymbol{\pi}}_n)$, because $\hat{\boldsymbol{\pi}}_n$ is a minimizer of $L_n(\boldsymbol{\pi})$ by definition. Therefore, by definition of $L_n(\boldsymbol{\pi})$, it suffices to show that for all (a, j) and $\varepsilon > 0$,

$$\Pr \left(\frac{1}{|I_a(\bar{\boldsymbol{\pi}}_n)|} \sum_{i \in I_a(\bar{\boldsymbol{\pi}}_n)} y_j^i - (\bar{\pi}_{aj})_n > \frac{\varepsilon}{md} \right) \rightarrow 0, \tag{EC.22}$$

where $I_a(\boldsymbol{\pi}) = \{i \in I | \mathbf{r}^i \in R_a(\boldsymbol{\pi})\}$. In the remainder of the proof, we suppress dependence of $\bar{\boldsymbol{\pi}}_n$ on n for conciseness. Next,

$$\Pr \left(\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I_a(\bar{\boldsymbol{\pi}})} y_j^i - \bar{\pi}_{aj} > \frac{\varepsilon}{md} \right) = \Pr \left(\sum_{i \in I_a(\bar{\boldsymbol{\pi}})} y_j^i - \bar{\pi}_{aj} | I_a(\bar{\boldsymbol{\pi}}) | > \frac{\varepsilon |I_a(\bar{\boldsymbol{\pi}})|}{md} \right),$$

$$\begin{aligned}
&\leq \Pr \left(\sum_{b \in A} \sum_{j \in J} \left| \sum_{i \in I_a(\bar{\pi})} y_j^i - \bar{\pi}_{aj} |I_a(\bar{\pi})| \right| > \frac{\varepsilon |I_a(\bar{\pi})|}{md} \right), \\
&= \Pr \left(n Z_n(\bar{\pi}) > \frac{\varepsilon |I_a(\bar{\pi})|}{md} \right), \\
&\leq \Pr \left(n Z_n(\boldsymbol{\pi}^0) > \frac{\varepsilon |I_a(\bar{\pi})|}{md} \right),
\end{aligned}$$

where the first line follows by multiplying both sides of the inequality by $|I_a(\bar{\pi})|$, the second line follows by non-negativity of the absolute value and summing over $b \in A$ and $j \in J$, the third line follows by definition of $Z_n(\boldsymbol{\pi})$, and the fourth line follows because $\bar{\pi} \in \operatorname{argmin}_{\boldsymbol{\pi} \in \Pi} Z_n(\boldsymbol{\pi})$, by definition. We can now write

$$\begin{aligned}
\Pr \left(n Z_n(\boldsymbol{\pi}^0) > \frac{\varepsilon |I_a(\bar{\pi})|}{md} \right) &= \Pr \left(\sum_{b \in A} \sum_{j \in J} \left| \sum_{i \in I_b(\boldsymbol{\pi}^0)} y_j^i - \pi_{bj}^0 |I_b(\boldsymbol{\pi}^0)| \right| > \frac{\varepsilon |I_a(\bar{\pi})|}{md} \right) \\
&\leq \sum_{b \in A} \sum_{j \in J} \Pr \left(\left| \sum_{i \in I_b(\boldsymbol{\pi}^0)} y_j^i - \pi_{bj}^0 |I_b(\boldsymbol{\pi}^0)| \right| > \frac{\varepsilon |I_a(\bar{\pi})|}{m^2 d^2} \right) \\
&\leq \sum_{b \in A} \sum_{j \in J} \sum_{k=0}^n \Pr \left(\left| \sum_{i \in I_b(\boldsymbol{\pi}^0)} y_j^i - \pi_{bj}^0 |I_b(\boldsymbol{\pi}^0)| \right| > \frac{\varepsilon k}{m^2 d^2} \mid |I_a(\bar{\pi})| = k \right) \Pr(|I_a(\bar{\pi})| = k) \\
&\leq 2 \sum_{b \in A} \sum_{j \in J} \sum_{k=0}^n \exp \left(- \left(\frac{\varepsilon k}{m^2 d^2} \right)^2 \cdot \frac{1}{|I_b(\boldsymbol{\pi}^0)|} \right) \Pr(|I_a(\bar{\pi})| = k) \\
&\leq 2d \sum_{b \in A} \sum_{k=0}^n \exp \left(- \left(\frac{\varepsilon k}{m^2 d^2} \right)^2 \cdot \frac{1}{n} \right) \Pr(|I_a(\bar{\pi})| = k),
\end{aligned}$$

where the first line follows by definition of $Z_n(\boldsymbol{\pi})$, the second line follows from the union bound, the third line follows by conditioning on $|I_a(\bar{\pi})|$, the fourth line follows by Hoeffding's inequality, and the fifth line follows by summing over $j \in J$ and because $n \geq |I_b(\boldsymbol{\pi}^0)|$ for all $b \in A$. Thus, to prove that (EC.22) holds, it remains to show that for each $a \in A$,

$$\sum_{k=0}^n \exp \left(- \left(\frac{\varepsilon k}{m^2 d^2} \right)^2 \cdot \frac{1}{n} \right) \Pr(I_a(\bar{\pi}) = k) \longrightarrow 0.$$

Next, using the fact that $\Pr(|I_a(\bar{\pi})| = k)$, $k = 1, \dots, n$ are binomial probabilities with parameter $\Pr(\mathbf{r} \in R_a(\bar{\pi}))$, and $\Pr(\mathbf{r} \in R_a(\boldsymbol{\pi})) > 0$ for all $a \in A$ and $\boldsymbol{\pi} \in \Pi$ (Assumption 1), it can be shown with some effort that for any $\delta \in (0, 1)$, there exists $n > 0$ and $K \in [0, n]$ such that $\Pr(|I_a(\boldsymbol{\pi})| \leq K) \leq \delta$ for any $\boldsymbol{\pi} \in \Pi$ and $\exp(-(\frac{\varepsilon k}{m^2 d^2})^2 \cdot \frac{1}{n}) < \delta$ for all $k \geq K$. It follows that

$$\sum_{k=0}^K \Pr(|I_a(\bar{\pi})| = k) = \Pr(|I_a(\bar{\pi})| \leq K) \leq \delta \tag{EC.23}$$

and

$$\sum_{k=K}^n \delta \cdot \Pr(|I_a(\bar{\pi})| = k) = \delta \sum_{k=K}^n \Pr(|I_a(\bar{\pi})| = k) \leq \delta. \quad (\text{EC.24})$$

Then we can write

$$\begin{aligned} \sum_{k=0}^n \exp\left(-\left(\frac{\varepsilon k}{m^2 d^2}\right)^2 \cdot \frac{1}{n}\right) \Pr(|I_a(\bar{\pi})| = k) &\leq \sum_{k=0}^{K-1} \Pr(|I_a(\bar{\pi})| = k) + \sum_{k=K}^n \delta \cdot \Pr(|I_a(\bar{\pi})| = k) \\ &\leq 2\delta, \end{aligned}$$

where the first term after the first inequality follows because $\exp(-(\frac{\varepsilon k}{m^2 d^2})^2 \cdot \frac{1}{n}) \leq 1$ for all $k \in [0, K-1]$ and $n \geq 0$, and the second term follows because $\exp(-(\frac{\varepsilon k}{m^2 d^2})^2 \cdot \frac{1}{n}) \leq \delta$ for all $k \geq K$. Letting $\delta \rightarrow 0$ yields the result. Part (iii). Because $|L_n(\pi_n^*) - L_n(\hat{\pi}_n)| \rightarrow 0$ by part (ii) above, and $|L_n(\hat{\pi}_n) - L_n(\pi^0)| \rightarrow 0$ by Theorem 1, $|L_n(\pi_n^*) - L_n(\pi^0)| \rightarrow 0$. The remainder of the proof follows by parallel argument to the proof of Theorem 1, with π_n^* in place of $\hat{\pi}_n$. \square

EC.5.2. Proof of Theorem 2

We first present two helpful supporting results in Lemmas EC.4 and EC.5. Lemma EC.4 is a concentration inequality that bounds the distance between an empirical mass function obtained from sampling from a discrete distribution and the discrete distribution itself. In Lemma EC.5, we define a Bernoulli random variable $e_a(\pi^0, \pi)$, whose value depends on the realization of \mathbf{r} , and is equal to 1 if the agent's optimal action is a under the true parameter π^0 but not under an alternative model π . Intuitively, the event $\{e_a(\pi^0, \pi) = 1\}$ represents a “mis-classification” of the agent action by the model π . Lemma EC.5 develops a bound on the probability that $e_a(\pi^0, \pi)$ is positive (i.e., equal to 1), which is the key result that we use to obtain the bound in Theorem 2.

LEMMA EC.4. *Let $\xi^1, \xi^2, \dots, \xi^n$ be i.i.d. discrete random variables with support $J = \{1, 2, \dots, d\}$ and mass function $\lambda_j = \Pr(\xi = j)$ for $j \in J$. Define $f_j = \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{\xi^i = j\}$ to be the empirical probability. Then*

$$\Pr\left(\sup_{j \in J} |f_j - \lambda_j| > \varepsilon\right) \leq 2 \exp(-n\varepsilon^2).$$

Proof. For convenience, let $\Lambda_j = \sum_{k=1}^j \lambda_k$ be the cumulative distribution and let $F_j = \sum_{k=1}^j f_k$ be the empirical cumulative distribution. With some effort, it can be shown that $\sup_{j \in J} |f_j - \lambda_j| \leq 2 \sup_{j \in J} |F_j - \Lambda_j|$. It follows that for any $\varepsilon > 0$,

$$\Pr\left(\sup_{j \in J} |f_j - \lambda_j| > \varepsilon\right) \leq \Pr\left(2 \sup_{j \in J} |F_j - \Lambda_j| > \varepsilon\right) = \Pr\left(\sup_{j \in J} |F_j - \Lambda_j| > \varepsilon/2\right) \leq 2 \exp(-n\varepsilon^2),$$

where the final inequality is the Dvoretzky–Kiefer–Wolfowitz inequality (Massart 1990). \square

LEMMA EC.5. *Let Assumption 3 hold. Let $e_a(\boldsymbol{\pi}^0, \boldsymbol{\pi})$ be a Bernoulli random variable equal to 1 if the events $\{\mathbf{r} \in R_a(\boldsymbol{\pi}^0)\}$ and $\{r \notin R_a(\boldsymbol{\pi})\}$ both occur. Then there exists $\bar{\boldsymbol{\pi}} \in \Pi$, and constants $\delta_1 \in (0, 1)$ and $\delta_2 \in (0, 1)$ such that*

$$\Pr(e_a(\boldsymbol{\pi}^0, \boldsymbol{\pi}) > 0) \leq 4m(1 - \delta_1(1 - \delta_2))^n \quad (\text{EC.25})$$

for all $a \in A$.

Proof. The proof proceeds in four steps. In the first step, we construct $\bar{\boldsymbol{\pi}}$. In the second and third steps, we prove two useful inequalities. In the fourth step, we prove the inequality (EC.25) in the lemma statement. Step 1. By Assumption 3, for each $a \in A$, there exists a cluster $s(a) \in S$ such that $B_{s(a)} \in R_a(\boldsymbol{\pi}^0)$. Pick $s(a)$ accordingly for each $a \in A$, and set $w_{as} = 1$ for $s = s(a)$ and $w_{as} = 0$ for $s \in S \setminus s(a)$. Then let $\bar{\boldsymbol{\pi}}_a = \sum_{s \in S} \mathbf{v}_s w_{as}$, for $a \in A$. Note that by construction, $\bar{\boldsymbol{\pi}}_a = \mathbf{v}_{s(a)}$ for $a \in A$. Step 2. First, we define the following useful quantity, which will be used throughout the proof of Theorem 2:

$$u_a(\mathbf{r}, \boldsymbol{\pi}) = \inf_{b \in A \setminus a} \left\{ \sum_{j \in J} (\pi_{aj} - \pi_{bj}) r_j + c_b - c_a \right\}.$$

Intuitively, $u_a(\mathbf{r}, \boldsymbol{\pi})$ represents the difference in agent utility between action a and the highest-utility action other than a . Based on this definition, note that an action a is optimal if and only if $u_a(\mathbf{r}, \boldsymbol{\pi}) \geq 0$. In this step, we show that there exists $\delta_1 \in (0, 1)$ and $\delta_2 \in (0, 1)$ such that $\Pr(\mathbf{r} \in B_s) \geq \delta_1$ for all $s \in S$, and

$$\int_{R_a(\boldsymbol{\pi}^0)} \exp\left(-k \left(\frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right)^2\right) f(\mathbf{r}) d\mathbf{r} \leq \delta_2 \quad (\text{EC.26})$$

for all $k \geq 0$ and $a \in A$. Recall that $\bar{r} = \sup_{\mathbf{r} \in R} \|\mathbf{r}\|_0 < \infty$; that is, \bar{r} is an upperbound on the largest agent payment r_j . First, for the existence of $\delta_1 \in (0, 1)$, it follows from the continuity of $f(\mathbf{r})$ on R (Assumption 1) and $B_s \subseteq R$ that $\Pr(\mathbf{r} \in B_s) > 0$ for all $s \in S$. Letting $\delta_1 = \inf_{s \in S} \Pr(\mathbf{r} \in B_s)$ implies $\Pr(\mathbf{r} \in B_s) \geq \delta_1$ for all $s \in S$, as desired. Next, for the existence of $\delta_2 \in (0, 1)$ such that (EC.26) holds, first consider the case where $k = 0$. Then

$$\int_{R_a(\boldsymbol{\pi}^0)} \exp\left(-0 \left(\frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right)^2\right) f(\mathbf{r}) d\mathbf{r} = \int_{R_a(\boldsymbol{\pi}^0)} f(\mathbf{r}) d\mathbf{r} < 1,$$

for all $a \in A$, where the final inequality follows because $R_a(\boldsymbol{\pi}^0) \subset R$ for all $a \in A$ (Assumption 1). It follows there exists $\tilde{\delta} \in (0, 1)$ such that (EC.26) holds for $k = 0$ and all $a \in A$. Now consider the case where $k \geq 1$. Observe that for all r such that $u_a(\mathbf{r}, \boldsymbol{\pi}^0) > 0$, $\exp(-(u_a(\mathbf{r}, \boldsymbol{\pi}^0)/(2\bar{r}d))^2) < 1$. Therefore,

for each $a \in A$ there exists $\tilde{\delta}_a \in (0, 1)$ such that if $u_a(\mathbf{r}, \boldsymbol{\pi}^0) > 0$, then $0 < \exp(-(u_a(\mathbf{r}, \boldsymbol{\pi}^0)/(2\bar{r}d))^2) \leq \tilde{\delta}_a < 1$. Next, exponentiating and integrating both sides of the preceding inequality yields

$$\int_{R_a(\boldsymbol{\pi}^0)} \exp\left(-k \left(\frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right)^2\right) f(\mathbf{r}) d\mathbf{r} \leq \int_{R_a(\boldsymbol{\pi}^0)} (\tilde{\delta}_a)^k f(\mathbf{r}) d\mathbf{r} \leq (\tilde{\delta}_a)^k.$$

The result follows from setting $\delta_2 = \sup\{\tilde{\delta}, \sup_{a \in A} \tilde{\delta}_a\}$. Step 3. In this step, we show that for $\bar{\pi}$ as constructed in step 1 and any $a \in A$,

$$\Pr\left(\sup_{j \in J} |\bar{\pi}_{bj} - \pi_{bj}^0| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right) \leq 2(1 - \delta_1(1 - \delta_2))^n$$

for all $b \in A$. For convenience, let $I_s = \{i \in I | \mathbf{r}^i \in B_s\}$. Then for each $b \in A$,

$$\begin{aligned} \Pr\left(\sup_{j \in J} |\bar{\pi}_{bj} - \pi_{bj}^0| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right) &= \Pr\left(\sup_{j \in J} \left|\frac{1}{|I_{s(a)}|} \sum_{i \in I_{s(a)}} y_j^i - \pi_{bj}^0\right| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right) \\ &= \sum_{k=0}^n \Pr\left(\sup_{j \in J} \left|\frac{1}{|I_{s(a)}|} \sum_{i \in I_{s(a)}} y_j^i - \pi_{bj}^0\right| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \middle| |I_{s(a)}| = k\right) \Pr(|I_{s(a)}| = k) \\ &= \sum_{k=0}^n \left[\int_R \Pr\left(\sup_{j \in J} \left|\frac{1}{k} \sum_{i \in I_{s(a)}} y_j^i - \pi_{bj}^0\right| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \middle| |I_{s(a)}| = k\right) f(\mathbf{r}) d\mathbf{r} \right] \Pr(|I_{s(a)}| = k) \\ &\leq 2 \sum_{k=0}^n \left[\int_{R_a(\boldsymbol{\pi}^0)} \exp\left(-k \left(\frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d}\right)^2\right) f(\mathbf{r}) d\mathbf{r} \right] \Pr(|I_{s(a)}| = k) \\ &\leq 2 \sum_{k=0}^n \tilde{\delta}^k \Pr(|I_{s(a)}| = k), \end{aligned}$$

The first line follows because $\bar{\pi}_{aj} = v_{s(a),j} = (1/|I_{s(a)}|) \sum_{i \in I_{s(a)}} y_j^i$ for all (a, j) , by construction of v_s (Algorithm 1) and $\bar{\pi}$ (Step 1). The second line follows from the total probability rule, because $|I_{s(a)}|$ is a binomial random variable. The third line follows from conditioning on r and integrating over R , and because y_j^i for $i \in |I_{s(a)}|$ are independent of \mathbf{r} . The fourth line follows from Lemma EC.4, and because $R_a(\boldsymbol{\pi}^0) \subseteq R$. The fifth line follows from Step 2 of the proof. Next, note that $|I_s|$ is the number of observations that are contained in the ball B_s . Therefore, $|I_s|$ is a binomial random variable with parameter $\lambda_s = \Pr(\mathbf{r} \in B_s)$. We can now write

$$\begin{aligned} 2 \sum_{k=0}^n \delta_2^k \cdot \Pr(|I_{s(a)}| = k) &= 2 \sum_{k=0}^n (\delta_2 \cdot \lambda_{s(a)})^k \frac{n!}{k!(n-k)!} (1 - \lambda_{s(a)})^{n-k} \\ &= 2(1 - \lambda_{s(a)}(1 - \delta_2))^n \\ &\leq 2(1 - \delta_1(1 - \delta_2))^n, \end{aligned}$$

where the first equality follows from writing the binomial probabilities $\Pr(|I_{s(a)}| = k)$ explicitly and grouping terms raised to the k^{th} power, the second equality follows immediately from the binomial

theorem, and the third inequality follows from Step 2. The result follows. Step 4. We now prove inequality (EC.25). Note that by definition, $e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > 0$ implies $r \notin R_a(\bar{\boldsymbol{\pi}})$, which implies there exists $b \in A$ such that $\sum_{j \in J} (\bar{\pi}_{bj} - \bar{\pi}_{aj})r_j > c_b - c_a$. Therefore, by the union bound,

$$\Pr(e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > 0) \leq \sum_{b \in A \setminus a} \Pr \left(\sum_{j \in J} (\bar{\pi}_{bj} - \bar{\pi}_{aj})r_j > c_b - c_a \right).$$

In the remainder of the proof, we bound $\Pr \left(\sum_{j \in J} (\bar{\pi}_{bj} - \bar{\pi}_{aj})r_j > c_b - c_a \right)$ for each $b \in A \setminus a$. For each $b \in A \setminus a$, we have

$$\begin{aligned} & \Pr \left(\sum_{j \in J} (\bar{\pi}_{bj} - \bar{\pi}_{aj})r_j > c_b - c_a \right), \\ & \leq \Pr \left(\sum_{j \in J} (\bar{\pi}_{bj} - \pi_{bj}^0)r_j - \sum_{j \in J} (\bar{\pi}_{aj} - \pi_{aj}^0)r_j > u_a(\mathbf{r}, \boldsymbol{\pi}^0) \right), \\ & \leq \Pr \left(\left\{ \sum_{j \in J} (\bar{\pi}_{bj} - \pi_{bj}^0)r_j > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2} \right\} \cup \left\{ \sum_{j \in J} (\pi_{aj}^0 - \bar{\pi}_{aj})r_j > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2} \right\} \right), \\ & \leq \Pr \left(\left\{ \sup_{j \in J} |\bar{\pi}_{bj} - \pi_{bj}^0| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \right\} \cup \left\{ \sup_{j \in J} |\pi_{aj}^0 - \bar{\pi}_{aj}| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \right\} \right), \\ & \leq \Pr \left(\sup_{j \in J} |\bar{\pi}_{bj} - \pi_{bj}^0| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \right) + \Pr \left(\sup_{j \in J} |\pi_{aj}^0 - \bar{\pi}_{aj}| > \frac{u_a(\mathbf{r}, \boldsymbol{\pi}^0)}{2\bar{r}d} \right), \\ & \leq 4(1 - \delta_1(1 - \delta_2))^n. \end{aligned}$$

The first inequality follows because $u_a(\mathbf{r}, \boldsymbol{\pi}^0) \leq \sum_{j \in J} (\pi_{aj}^0 - \pi_{bj}^0)r_j + c_b - c_a$ for all $b \in A$, by definition of $u_a(\mathbf{r}, \boldsymbol{\pi}^0)$. The second inequality follows because $\sum_{j \in J} (\pi_{bj} - \pi_{bj}^0)r_j - \sum_{j \in J} (\bar{\pi}_{aj} - \pi_{aj}^0)r_j > u_a(\mathbf{r}, \boldsymbol{\pi}^0)$ implies at least one of $\sum_{j \in J} (\bar{\pi}_{bj} - \pi_{bj}^0)r_j > u_a(\mathbf{r}, \boldsymbol{\pi}^0)/2$ or $\sum_{j \in J} (\pi_{aj}^0 - \bar{\pi}_{aj})r_j > u_a(\mathbf{r}, \boldsymbol{\pi}^0)/2$ holds. The third inequality follows because $|J| \sup_{j \in J} \{(\pi_{aj}^0 - \bar{\pi}_{aj})r_j\} \geq \sum_{j \in J} (\pi_{aj}^0 - \bar{\pi}_{aj})r_j$, $|J| = d$, and $\bar{r} \geq r_j$ for all $j \in J$. The fourth inequality follows from the union bound. The fifth inequality follows from Step 3. The result follows by summing the final inequality over $b \in A$. \square

Proof of Theorem 2. We prove the result by constructing a sequence of feasible solutions $\bar{\boldsymbol{\pi}}_n$ to PA-D, and bounding the objective function under $\bar{\boldsymbol{\pi}}_n$. For conciseness we suppress dependence of $\bar{\boldsymbol{\pi}}_n$ and $\tilde{\boldsymbol{\pi}}_n$ on n . By Assumption 3, for each $a \in A$, there exists a cluster $s(a) \in S$ such that $B_{s(a)} \in R_a(\boldsymbol{\pi}^0)$. Pick $s(a)$ accordingly for each $a \in A$. Set $\bar{w}_{as} = 1$ for $s = s(a)$ and $\bar{w}_{as} = 0$ for all $s \in S \setminus s(a)$. Fix $\mathbf{w} = \bar{\mathbf{w}}$ in PA-D and let $(\bar{\mathbf{z}}, \bar{\mathbf{x}}, \bar{\boldsymbol{\phi}})$ be the solution of the resulting subproblem. Then $(\bar{\mathbf{w}}, \bar{\mathbf{z}}, \bar{\mathbf{x}}, \bar{\boldsymbol{\phi}})$ is a feasible solution to PA-D, and the associated estimate is given by

$$\bar{\boldsymbol{\pi}}_a = \sum_{s \in S} \mathbf{v}_s \bar{w}_{as} = \mathbf{v}_{s(a)}, \quad a \in A.$$

Because $\bar{\boldsymbol{\pi}}$ is attained at a feasible solution to PA-D, $Z_n(\hat{\boldsymbol{\pi}}) \leq Z_n(\bar{\boldsymbol{\pi}})$. It follows that

$$\Pr(|Z_n(\hat{\boldsymbol{\pi}}) - Z_n(\tilde{\boldsymbol{\pi}})| > \varepsilon) \leq \Pr(Z_n(\tilde{\boldsymbol{\pi}}) > \varepsilon) \leq \Pr(Z_n(\bar{\boldsymbol{\pi}}) > \varepsilon),$$

for any $\varepsilon > 0$. Therefore, it suffices to show that $\Pr(Z_n(\bar{\boldsymbol{\pi}}) > \varepsilon) = O(n^2\kappa^n)$ for some $\kappa \in (0, 1)$.

Because $Z_n(\boldsymbol{\pi}) = \sum_{a \in A} \sum_{j \in J} \frac{|I_a(\boldsymbol{\pi})|}{n} |\pi_{aj} - \omega_{aj}(\bar{\boldsymbol{\pi}})|$ by definition, we shall bound $Z_n(\bar{\boldsymbol{\pi}})$ by bounding $\frac{|I_a(\boldsymbol{\pi})|}{n} |\bar{\pi}_{aj} - \omega_{aj}(\bar{\boldsymbol{\pi}})|$ for each (a, j) . Next, pick any (a, j) . Then

$$\begin{aligned} \frac{|I_a(\bar{\boldsymbol{\pi}})|}{n} |\bar{\pi}_{aj} - \omega_{aj}(\bar{\boldsymbol{\pi}})| &= \frac{|I_a(\bar{\boldsymbol{\pi}})|}{n} |\bar{\pi}_{aj} - \pi_{aj}^0 + \pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0) + \omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})|, \\ &\leq \frac{|I_a(\bar{\boldsymbol{\pi}})|}{n} (|\bar{\pi}_{aj} - \pi_{aj}^0| + |\pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0)| + |\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})|), \\ &\leq |\bar{\pi}_{aj} - \pi_{aj}^0| + |\pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0)| + |\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})|, \end{aligned} \quad (\text{EC.27})$$

where the second line follows by the triangle inequality, and the third line follows because $|I_a(\boldsymbol{\pi})| \leq n$ for any $\boldsymbol{\pi} \in \Pi$. The remainder of the proof proceeds in three steps. In each step, we bound one of the terms in the right-hand side of (EC.27), from left to right. Step 1. For the first term, $|\bar{\pi}_{aj} - \pi_{aj}^0|$, note

$$\begin{aligned} \Pr(|\bar{\pi}_{aj} - \pi_{aj}^0| > \varepsilon) &= \sum_{k=0}^n \Pr(|\bar{\pi}_{aj} - \pi_{aj}^0| > \varepsilon | |I_{s(a)}| = k) \Pr(|I_{s(a)}| = k) \\ &= \sum_{k=0}^n \Pr\left(\left|\frac{1}{|I_{s(a)}|} \sum_{i \in I_{s(a)}} y_j^i - \pi_{aj}^0\right| > \varepsilon \mid |I_{s(a)}| = k\right) \Pr(|I_{s(a)}| = k) \\ &\leq \sum_{k=0}^n 2 \exp(-\varepsilon^2 k) \Pr(|I_{s(a)}| = k). \end{aligned}$$

The first line follows by conditioning on $|I_{s(a)}| = k$. The second line follows because $\bar{\pi}_{aj} = v_{s(a),j} = (1/|I_{s(a)}|) \sum_{i \in I_{s(a)}} y_j^i$ by the construction of $\bar{\pi}_{aj}$ above and the definition of $\mathbf{v}_{s(a)}$ (Algorithm 1). For the third line, note that $(1/|I_{s(a)}|) \sum_{i \in I_{s(a)}} y_j^i$ is the empirical mass function for the independent variables ξ^i for all $i \in I_{s(a)}$, generated by $\boldsymbol{\pi}^0$. The inequality, therefore, follows by an application of Lemma EC.4. Next, because the \mathbf{r}^i are i.i.d. (Assumption 1), $|I_{s(a)}|$ is a binomial random variable with parameter $\Pr(\mathbf{r} \in B_{s(a)})$. For conciseness, define $\lambda_{s(a)} = \Pr(\mathbf{r} \in B_{s(a)})$. We now have

$$\begin{aligned} \sum_{k=0}^n (2 \exp(-\varepsilon^2 k)) \Pr(|I_{s(a)}| = k) &= \sum_{k=0}^n (2 \exp(-\varepsilon^2 k)) \frac{n!}{k!(n-k)!} \lambda_{s(a)}^k (1 - \lambda_{s(a)})^{n-k} \\ &= 2 \sum_{k=0}^n (\lambda_{s(a)} \exp(-\varepsilon^2))^k \frac{n!}{k!(n-k)!} (1 - \lambda_{s(a)})^{n-k}, \end{aligned}$$

where the first equality follows from writing the binomial probabilities $\Pr(|I_{s(a)}| = k)$ explicitly, and the second equality follows from grouping terms raised to the k^{th} power. Finally, we have

$$\begin{aligned} 2 \sum_{k=0}^n (\lambda_{s(a)} \exp(-\varepsilon^2))^k \frac{n!}{k!(n-k)!} (1 - \lambda_{s(a)})^{n-k} &= 2(1 - \lambda_{s(a)}(1 - \exp(-\varepsilon^2)))^n \\ &\leq 2(1 - \delta_1(1 - \exp(-\varepsilon^2)))^n, \end{aligned}$$

for some $\delta_1 \in (0, 1)$, where the equality follows by applying the binomial theorem, and the inequality follows by letting $\delta_1 = \inf_{s \in S} \lambda_{s(a)}$. Therefore,

$$\Pr(|(\bar{\pi}_{aj} - \pi_{aj}^0) > \varepsilon| \leq 2(1 - \delta_1(1 - \exp(-\varepsilon^2)))^n. \quad (\text{EC.28})$$

Step 2. Next, we bound the second term $|\pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0)|$ in (EC.27). Because $\omega_{aj}(\boldsymbol{\pi}^0) = (1/|I_a(\boldsymbol{\pi}^0)|) \sum_{i \in I_a(\boldsymbol{\pi}^0)} y_j^i$, it follows immediately from Lemma EC.4 that

$$\Pr(|\pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0)| > \varepsilon) \leq 2 \exp(-n\varepsilon^2). \quad (\text{EC.29})$$

Step 3. Next, we bound the third term $|\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})|$ in (EC.27). For convenience, let $\beta_{aj}(\boldsymbol{\pi})$ be a Bernoulli random variable equal to 1 if the events $\{i \in I_a(\boldsymbol{\pi})\}$ and $\{y_j = 1\}$ both occur, and let $\beta_{aj}^i(\boldsymbol{\pi})$ be the realized value of $\beta_{aj}(\boldsymbol{\pi})$ in the i^{th} observation. Using the definition of β_{aj}^i , it follows that

$$\begin{aligned} |\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})| &= \left| \frac{1}{|I_a(\boldsymbol{\pi}^0)|} \sum_{i \in I_a(\boldsymbol{\pi}^0)} y_j^i - \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I_a(\bar{\boldsymbol{\pi}})} y_j^i \right| \\ &= \left| \frac{1}{|I_a(\boldsymbol{\pi}^0)|} \sum_{i \in I} \beta_{aj}^i(\boldsymbol{\pi}^0) - \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) \right|. \end{aligned}$$

Next, we have

$$\begin{aligned} \left| \frac{1}{|I_a(\boldsymbol{\pi}^0)|} \sum_{i \in I} \beta_{aj}^i(\boldsymbol{\pi}^0) - \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) \right| &\leq \left| \max \left\{ \frac{1}{|I_a(\boldsymbol{\pi}^0)|}, \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \right\} \sum_{i \in I} \beta_{aj}^i(\boldsymbol{\pi}^0) - \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) \right| \\ &\leq \left| \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}})) \right| + \left| \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|} \right| \left| \sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) \right| \\ &\leq \left| \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}})) \right| + \left| \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|} \right| \left| \sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) \right|, \end{aligned} \quad (\text{EC.30})$$

where the first two lines are consequences of the $\max\{\cdot\}$ operator, and the third line follows from the triangle inequality. We shall bound each of the two terms on the right-hand side of (EC.30) separately. For the first term on the right-hand side of (EC.30), observe that by definition of $\beta_{aj}^i(\boldsymbol{\pi})$, $\sum_{j \in J} \beta_{aj}^i(\boldsymbol{\pi}) = 1$ if and only if $i \in I_a(\boldsymbol{\pi})$. Recall from Lemma EC.5 that $e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}})$ is a Bernoulli variable equal to 1 if the events $\{\mathbf{r} \in R_a(\boldsymbol{\pi}^0)\}$ and $\{r \notin R_a(\bar{\boldsymbol{\pi}})\}$ both occur. Further, note $\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}}) > 0$ implies $\mathbf{r}^i \notin R_a(\bar{\boldsymbol{\pi}})$, which implies $e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) = 1$. Therefore,

$$\left| \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}})) \right| \leq \frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}).$$

Next, for any $\varepsilon \in (0, 1)$,

$$\begin{aligned} \Pr \left(\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon \right) &\leq \Pr \left(\sum_{i \in I} e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon \right) \\ &\leq \sum_{i \in I} \Pr(e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon) \\ &\leq n \Pr(e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon), \end{aligned}$$

where the first inequality follows because $|I_a(\bar{\boldsymbol{\pi}})| \geq 1$ by Assumption 3 and the construction of $\bar{\boldsymbol{\pi}}$, the second inequality follows from applying the union bound, and the third inequality follows because the e_a^i are i.i.d. (Assumption 1). Next, by Lemma EC.5, $\Pr(e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon) \leq 4m(1 - \delta_1(1 - \delta_2))^n$. Therefore, for any $\varepsilon \in (0, 1)$,

$$\Pr\left(\left|\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} \sum_{i \in I} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}}))\right| > \varepsilon\right) \leq 4mn(1 - \delta_1(1 - \delta_2))^n. \quad (\text{EC.31})$$

Next, we bound the second term in (EC.30). For any $\varepsilon \in (0, 1)$, note

$$\Pr\left(\left|\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|}\right| \left|\sum_{i \in I} \beta_{aj}^i(\bar{\boldsymbol{\pi}}) > \varepsilon\right.\right) \leq n \Pr\left(\left|\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|}\right| > \varepsilon\right),$$

which follows from the union bound and because $\beta_{aj}^i(\bar{\boldsymbol{\pi}}) \leq 1$ for all $i \in I$. Next, note that

$$\left|\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|}\right| \leq \left||I_a(\boldsymbol{\pi}^0)| - |I_a(\bar{\boldsymbol{\pi}})|\right|,$$

which follows because $|I_a(\boldsymbol{\pi}^0)| \geq 1$ and $|I_a(\bar{\boldsymbol{\pi}})| \geq 1$ by Assumption 3 and by construction of $\bar{\boldsymbol{\pi}}$. Therefore,

$$\begin{aligned} n \cdot \Pr\left(\left|\frac{1}{|I_a(\bar{\boldsymbol{\pi}})|} - \frac{1}{|I_a(\boldsymbol{\pi}^0)|}\right| > \varepsilon\right) &\leq n \cdot \Pr\left(\left||I_a(\boldsymbol{\pi}^0)| - |I_a(\bar{\boldsymbol{\pi}})|\right| > \varepsilon\right) \\ &= n \cdot \Pr\left(\left|\sum_{i \in I} \sum_{j \in J} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}}))\right| > \varepsilon\right) \\ &\leq n \cdot \Pr\left(\sum_{i \in I} e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon\right) \\ &\leq n^2 \cdot \Pr(e_a(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}}) > \varepsilon) \\ &\leq 4mn^2(1 - \delta_1(1 - \delta_2))^n, \end{aligned} \quad (\text{EC.32})$$

where the second line follows because $|I_a(\boldsymbol{\pi})| = \sum_{i \in I} \sum_{j \in J} \beta_{aj}^i(\boldsymbol{\pi})$ by definition of $\beta_{aj}^i(\boldsymbol{\pi})$, the third line follows because $\sum_{j \in J} (\beta_{aj}^i(\boldsymbol{\pi}^0) - \beta_{aj}^i(\bar{\boldsymbol{\pi}})) \leq e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}})$ for all $i \in I$, the fourth line follows by applying the union bound and noting that the $e_a^i(\boldsymbol{\pi}^0, \bar{\boldsymbol{\pi}})$ are i.i.d. (Assumption 1), and the fifth line follows from Lemma EC.5. Therefore, combining (EC.31) and (EC.32) produces the following bound on the third term in (EC.27) for any $\varepsilon \in (0, 1)$:

$$\Pr(|\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})| > \varepsilon) \leq 4mn(n+1)(1 - \delta_1(1 - \delta_2))^n. \quad (\text{EC.33})$$

Combining (EC.27), (EC.28), (EC.29) and (EC.33) and applying the union bound yields

$$\begin{aligned} &\Pr\left(\frac{|I_a(\bar{\boldsymbol{\pi}})|}{n} |\bar{\pi}_{aj} - \omega_{aj}(\bar{\boldsymbol{\pi}})| > \varepsilon\right) \\ &\leq \Pr(|\bar{\pi}_{aj} - \pi_{aj}^0| > \varepsilon/3) + \Pr(|\pi_{aj}^0 - \omega_{aj}(\boldsymbol{\pi}^0)| > \varepsilon/3) + \Pr(|\omega_{aj}(\boldsymbol{\pi}^0) - \omega_{aj}(\bar{\boldsymbol{\pi}})| > \varepsilon/3) \\ &\leq 2\left(\exp(-n(\varepsilon/3)^2) + (1 - \delta_1(1 - \exp(-(\varepsilon/3)^2)))^n\right) + 2mn(n+1)(1 - \delta_1(1 - \delta_2))^n. \end{aligned} \quad (\text{EC.34})$$

for any $\varepsilon \in (0, 1)$ and each $a \in A$ and $j \in J$. It follows that for any $\varepsilon \in (0, 1)$,

$$\Pr(Z_n(\bar{\pi}) > \varepsilon) = \Pr\left(\sum_{a \in A} \sum_{j \in J} \frac{|I_a(\bar{\pi})|}{n} |\bar{\pi}_{aj} - \omega_{aj}(\bar{\pi})| > \varepsilon\right) \quad (\text{EC.35})$$

$$\leq \sum_{a \in A} \sum_{j \in J} \Pr\left(|\bar{\pi}_{aj} - \omega_{aj}(\bar{\pi})| > \frac{\varepsilon}{md}\right) \\ \leq 2md(\exp(-n(\varepsilon/(3md))^2) + (1 - \delta_1(1 - \exp(-(\varepsilon/(3md))^2)))^n) \quad (\text{EC.36})$$

$$+ 2mn(n+1)(1 - \delta_1(1 - \delta_2))^n). \quad (\text{EC.37})$$

where the first line follows by definition of $Z_n(\pi)$, the second line follows from applying the union bound over all $a \in A$ and $j \in J$ and noting $m = |A|$ and $d = |J|$, and the third line follows from (EC.34). Note that the third term in (EC.37) is dominant, which implies $\Pr(Z_n(\bar{\pi}) > \varepsilon) \leq O(n^2 \kappa^n)$, where $\kappa = 1 - \delta_1(1 - \delta_2)$. Lastly, note $\kappa \in (0, 1)$ because $\delta_1 \in (0, 1)$ and $\delta_2 \in (0, 1)$. \square

EC.5.3. Proofs of Proposition 2, Theorem 3, and Theorem 4

Proof of Proposition 2. The proof proceeds in two steps. First, we show $Z_n(\tilde{\pi}_n) \rightarrow 0$. Second, we prove the main result. Step 1. Note that for any $\bar{\pi}$ attained at a feasible solution to PA-D(S), $0 \leq Z_n(\tilde{\pi}_n) \leq Z_n(\bar{\pi})$. Therefore, we prove the result by constructing a feasible solution $(\bar{\pi}, \bar{x}, \bar{w}, \bar{z}, \bar{\phi})$ for each $n \geq 0$ and showing $Z_n(\bar{\pi}) \rightarrow 0$. Because $Z_n(\pi^0) \rightarrow 0$ by Lemma EC.3, it suffices to show $\bar{\pi} \rightarrow \pi_0$. By Assumption 3, for each $a \in A$ there exists $s \in S$ such that $B_s \subseteq R_a(\pi^0)$. For each $n \geq 0$, let $(\bar{\pi}, \bar{x}, \bar{w}, \bar{z}, \bar{\phi})$ be constructed as follows: for each $a \in A$, set $\bar{w}_{as} = 1$ for $s = s(a)$ and $\bar{w}_{as} = 0$ for all $s \in S \setminus s(a)$. Fix $w = \bar{w}$ in PA-D and let $(\bar{\pi}, \bar{x}, \bar{z}, \bar{\phi})$ be the solution of the resulting subproblem. Then

$$\Pr\left(\sup_{j \in J} |\bar{\pi}_{aj} - \pi_{aj}^0| > \varepsilon\right) \leq \Pr\left(\sup_{j \in J} |(v_{s(a)j} - \pi_{aj}^0)| > \varepsilon\right) \\ \leq \Pr\left(\sup_{j \in J} \left|\frac{1}{|I_{s(a)}|} \sum_{i \in I_{s(a)}} y_j^i - \pi_{aj}^0\right| > \varepsilon\right) \\ \leq 2 \exp(-\varepsilon^2 |I_{s(a)}|),$$

where the first line follows from constraint (13i), the second line follows by definition of $\mathbf{v}_{s(a)}$, and the third line follows from Lemma EC.4. Because $f(\mathbf{r})$ is continuous on R (Assumption 1), $|I_{s(a)}| \rightarrow \infty$ as $n \rightarrow \infty$ for all $a \in A$. Therefore, $\bar{\pi}_n \rightarrow \pi_0$, as desired. Step 2. Note $0 \leq Z_n(\pi_n^*) \leq Z_n(\tilde{\pi}_n)$ by definition of π_n^* . Because $Z_n(\tilde{\pi}_n) \rightarrow 0$ by Step 1, it follows that $Z_n(\pi_n^*) \rightarrow 0$. Therefore, $|Z_n(\pi_n^*) - Z_n(\tilde{\pi}_n)| \rightarrow 0$, as desired. \square

Proof of Theorem 3. Let \tilde{w} be obtained at an optimal solution to PA-D, and define $\tilde{V} = \{\mathbf{v}_s \in V \mid \sum_{a \in A} \tilde{w}_{as} = 1\}$. Let \tilde{S} index the candidate distributions in \tilde{V} . Let V_T^+ denote the

candidate distributions at termination of Algorithm 2, with index set S_T^+ . For conciseness, we suppress dependence of \tilde{V} , V_T^+ , \tilde{S} and S_T^+ on n . The proof proceeds in two steps. First, we show $|Z_n(\tilde{\pi}_n) - Z_n(\pi_n^+)| \rightarrow 0$. Second, we prove the main result. Step 1. Observe that if $\tilde{V} \subseteq V_T^+$, then $Z_n(\pi_n^+) \leq Z_n(\tilde{\pi}_n)$. By optimality of $\tilde{\pi}_n$ with respect to PA-D(S), $Z_n(\tilde{\pi}_n) \leq Z_n(\pi_n^+)$. Thus $|Z_n(\tilde{\pi}_n) - Z_n(\pi_n^+)| \rightarrow 0$ if $\tilde{V} \subseteq V_T^+$. It therefore suffices to show that for any $\mathbf{v}_s \in \tilde{V}$, there exists $\mathbf{v}_{s'} \in V_T^+$ such that $|\mathbf{v}_s - \mathbf{v}_{s'}| \rightarrow 0$. Suppose otherwise. Then for all $n \geq 0$, there exists $\mathbf{v}_s \in \tilde{V}$ such that $\mathbf{v}_s \neq \mathbf{v}_{s'}$ for all $\mathbf{v}_{s'} \in V_T^+$. By the weak law of large numbers, for each $s \in S$ there exists $\boldsymbol{\nu}_s \in \mathbb{R}^d$ such that $\mathbf{v}_s \rightarrow \boldsymbol{\nu}_s$ as $n_s \rightarrow \infty$. It follows that for all $n \geq 0$, there exists $s \in \tilde{S}$ such that $\boldsymbol{\nu}_s \neq \boldsymbol{\nu}_{s'}$ for all $s' \in S_T^+$. Note $v_{sj} = \psi_{sj}/(\sum_{j \in J} \psi_{sj})$. Then for all $n \geq 0$, there exists $s \in \tilde{S}$ such that $H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) < 0$ and $\boldsymbol{\nu}_s \neq \boldsymbol{\nu}_{s'}$ for all $s' \in S_T^+$. Next, it is straightforward to verify that because the density function $f(\mathbf{r})$ is continuous on R (Assumption 1), $n_s \rightarrow \infty$ and $n_{s'} \rightarrow \infty$ as $n \rightarrow \infty$. Hence, $\Pr(H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0 | \boldsymbol{\nu}_s \neq \boldsymbol{\nu}_{s'}) \rightarrow 1$, a contradiction. Therefore, $|Z_n(\tilde{\pi}_n) - Z_n(\pi_n^+)| \rightarrow 0$. Step 2. By Step 1 of the proof of Proposition 2, we have $Z_n(\tilde{\pi}_n) \rightarrow 0$. Because $|Z_n(\tilde{\pi}_n) - Z_n(\pi_n^+)| \rightarrow 0$ from the first step of this proof, it follows that $Z_n(\pi_n^+) \rightarrow 0$. Following a parallel argument to the proof of Proposition 1, it can be shown that $Z_n(\pi_n^+) \rightarrow 0$ implies $|L_n(\pi_n^+) - L_n(\hat{\pi}_n)| \rightarrow 0$. Because $|L_n(\hat{\pi}_n) - L_n(\pi^0)| \rightarrow 0$ from the first step of the proof of Theorem 1, $|L_n(\pi_n^+) - L_n(\hat{\pi}_n)| \rightarrow 0$ implies $|L_n(\pi_n^+) - L_n(\pi^0)| \rightarrow 0$. The remainder of the proof follows by a parallel argument to the second step of the proof of Theorem 1, with π_n^+ in place of $\hat{\pi}_n$. \square

Proof of Theorem 4. The proof proceeds in two steps. First, we show $\Pr(T > m) \leq \alpha m S$. Second, we prove the main result. Step 1. By assumption, for each $s \in S$, there exists $a \in A$ such that $B_s \subseteq R_a(\pi^0)$. In words, this means that for each ball B_s , every contract $\mathbf{r} \in B_s$ induces the same action from the agent. Accordingly, each candidate distribution $s \in S$ is *mapped* to exactly one action $a \in A$. If there exists an $s \in S^+$ such that s is mapped to an action a , we say that the action a is *represented* in S^+ . For each $s \in S^-$, let E_s be the event that there exists $s'(s) \in S^+$ such that $B_s \subseteq R_a(\pi^0)$ and $B_{s'} \subseteq R_a(\pi^0)$; that is, E_s is the event that the action a that $s \in S^-$ maps to is already represented in S^+ . Next, pick an iteration t , and let s_t^* be the cluster selected in the t^{th} iteration of Algorithm 2. Then $E_{s_t^*}$ is the event that the candidate distribution selected in iteration t is mapped to an action that is already represented in S^+ . We can now write

$$\begin{aligned} \Pr(E_{s_t^*}) &\leq \Pr(\cup_{s \in S^-} \{ \cap_{s' \in S^+} \{ H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0 \}, E_s \}) \\ &\leq \sum_{s \in S^-} \Pr(\cap_{s' \in S^+} \{ H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0 \}, E_s) \\ &\leq \sum_{s \in S^-} \Pr\left(\cap_{s' \in S^+} \{ H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0 \} \middle| E_s\right) \Pr(E_s) \end{aligned}$$

$$\leq \sum_{s \in S^-} \Pr \left(\bigcap_{s' \in S^+} \{H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0\} \middle| E_s \right), \quad (\text{EC.38})$$

The first inequality follows by definition of s^τ (Algorithm 2) and E_s . The second inequality follows from the union bound, the third inequality follows from conditioning on E_s , and the fourth inequality follows because $\Pr(E_s) \leq 1$ for all $s \in S$. Next, for each $s \in S$, we have

$$\Pr \left(\bigcap_{s' \in S^+} \{H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0\} \middle| E_s \right) \leq \Pr \left(H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'(s)}) \geq 0 \middle| E_s \right) \quad (\text{EC.39})$$

$$\leq \Pr \left(H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'(s)}) \geq 0 \middle| \nu_s = \nu_{s'(s)} \right) \quad (\text{EC.40})$$

$$\leq \alpha, \quad (\text{EC.41})$$

where the first inequality follows from dropping the events $\{H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'}) \geq 0\}$ from the intersection for all elements in $S^+ \setminus s'(s)$, the second inequality follows because the event $\nu_s = \nu_{s'(s)}$ implies the event E_s , and the third inequality follows by definition of $H_\alpha(\boldsymbol{\psi}_s, \boldsymbol{\psi}_{s'})$. By combining (EC.38) and (EC.39) and noting $|S^-| \leq |S|$ for all $t = 1, \dots, T$, it follows that $\Pr(E_{s_t^*}) \leq \alpha|S|$. Next, note that the event $\{T > m\}$ implies $\{\cup_{t=0}^{m-1} E_{s_t^*}\}$ by the pigeonhole principle. Therefore, by the union bound,

$$\begin{aligned} \Pr(T > m) &\leq \Pr \left(\cup_{k=0}^{m-1} E_{s_k^*} \right) \\ &\leq \sum_{k=0}^{m-1} \Pr(E_{s_k^*}) \\ &\leq \sum_{k=0}^{m-1} \alpha|S| \\ &\leq \alpha m |S|, \end{aligned}$$

as desired. Step 2. Note $\mathbb{E}[T] = \sum_{k=0}^{\infty} \Pr(T > k)$ by definition of the expectation. It follows that

$$\begin{aligned} \sum_{k=0}^{\infty} \Pr(T > k) &= \sum_{k=0}^{|S|-1} \Pr(T > k) \\ &= \sum_{k=0}^{m-1} \Pr(T > k) + \sum_{k=m}^{|S|-1} \Pr(T > k) \\ &\leq m + \sum_{k=m}^{|S|-1} \Pr(T > k), \end{aligned}$$

where the first equality follows because T is bounded above by $|S|$, the second equality follows from separating the summation, and the inequality follows because $\sum_{k=0}^{m-1} \Pr(T > k) \leq m$. Next, note that the event $T > k$ implies $T > m$ for $k = m, \dots, |S| - 1$, and thus $\Pr(T > k) \leq \Pr(T > m)$ for $k = m, \dots, |S| - 1$. Therefore,

$$m + \sum_{k=m}^{|S|-1} \Pr(T > k) \leq m + \sum_{k=m}^{|S|-1} \Pr(T > m) = m + \Pr(T > m)(|S| - m) \leq m[1 + \alpha \cdot |S| \cdot (|S| - m)],$$

where the final inequality follows because $\Pr(T > m) \leq \alpha m |S|$ as established in Step 1. \square

EC.5.4. Proofs for Sections EC.3 and EC.4

Proof of Corollary EC.1. The proof proceeds similarly to the proof of Theorem 1: First, we show $|\sum_{k \in K} (L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0))| \rightarrow 0$. Second, we show $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n = \boldsymbol{\pi}^0$ if and only if Assumption EC.2 holds. Step 1. By Lemma EC.2, $L_n^k(\boldsymbol{\pi})$ is lower semicontinuous for all $k \in K$. Because the sum of lower semicontinuous functions is lower semicontinuous (Dietze and Schäuble 1985), it follows that $\sum_{k \in K} L_n^k(\boldsymbol{\pi})$ is lower semicontinuous. Next, we show $|\sum_{k \in K} (L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0))| \rightarrow 0$. By the triangle inequality,

$$\left| \sum_{k \in K} (L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0)) \right| \leq \sum_{k \in K} |L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0)|.$$

By the proof of Theorem 1 and Assumption EC.1, we have $|L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0)| \rightarrow 0$. Because $L_n^k(\boldsymbol{\pi}) \geq 0$ for any $\boldsymbol{\pi}$, it follows that $|\sum_{k \in K} (L_n^k(\hat{\boldsymbol{\pi}}_n) - L_n^k(\boldsymbol{\pi}^0))| \rightarrow 0$. Step 2. For each $k \in K$, define $L^k(\boldsymbol{\pi}) = \sum_{a \in A} \sum_{j \in J} |\pi_{aj}^k - \lim_{n \rightarrow \infty} \omega_{aj}^k(\boldsymbol{\pi})|$. Because $\sum_{k \in K} L_n^k(\boldsymbol{\pi})$ is lower semicontinuous, $|\sum_{k \in K} L_n^k(\hat{\boldsymbol{\pi}}_n) - \sum_{k \in K} L_n^k(\boldsymbol{\pi}^0)| \rightarrow 0$, and Π is compact, by Theorem 5.14 of Van der Vaart (2000),

$$\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n \in \underset{\boldsymbol{\pi} \in \Pi}{\text{argmin}} \sum_{k \in K} L^k(\boldsymbol{\pi}).$$

Next, by parallel argument to Lemma EC.3, it can be shown that $\boldsymbol{\pi}^0$ is the unique minimizer of $\sum_{k \in K} L^k(\boldsymbol{\pi})$ if and only if Assumption EC.2 holds. Suppose Assumption EC.2 holds. It follows that $L^k(\boldsymbol{\pi}^0) = 0$ and $L^k(\boldsymbol{\pi}) > 0$ for $\boldsymbol{\pi} \neq \boldsymbol{\pi}^0$. Therefore, $\underset{\boldsymbol{\pi} \in \Pi}{\text{argmin}} \sum_{k \in K} L_n^k(\boldsymbol{\pi}) = \boldsymbol{\pi}^0$, which implies $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n = \boldsymbol{\pi}^0$. Conversely, if Assumption EC.2 does not hold, then there exists $\tilde{\boldsymbol{\pi}} \neq \boldsymbol{\pi}^0$ such that $\sum_{k \in K} L(\tilde{\boldsymbol{\pi}}) = \sum_{k \in K} L(\boldsymbol{\pi}^0) = 0$, which implies $\text{plim}_{n \rightarrow \infty} \hat{\boldsymbol{\pi}}_n \neq \boldsymbol{\pi}^0$. \square

Proof of Proposition EC.1. For convenience, let $Z^t(\mathbf{w}^t, \boldsymbol{\varepsilon}^t)$ be the objective function value of PA-T under $(\mathbf{w}^t, \boldsymbol{\varepsilon}^t)$, and let $(\bar{\mathbf{w}}^t, \bar{\boldsymbol{\varepsilon}}^t)$ be the optimal solution to PA-T in iteration t of Algorithm 5. Let “ \rightarrow ” denote convergence in probability as $T \rightarrow \infty$. The proof proceeds in four steps. We prove useful supporting results in the first three steps: In Step 1 we show $\mathbf{v}_s^T \rightarrow \boldsymbol{\pi}_{a(\mathbf{r}^s)}^0$ for each $s \in S$; in Step 2 we show $Z^T(\bar{\mathbf{w}}^T, \bar{\boldsymbol{\varepsilon}}^T) \rightarrow 0$; in Step 3 we show $\Pr(\bar{w}_{as}^T = 1) \rightarrow 1$ for $a = a(\mathbf{r}^s)$ and $\Pr(\bar{w}_{as}^T = 0) \rightarrow 1$ for $a \neq a(\mathbf{r}^s)$, where $a(\mathbf{r}^s)$ is the agent’s true optimal action under \mathbf{r}^s and $\boldsymbol{\pi}^0$. In Step 4, we combine these results to prove the statement of Proposition EC.1. Step 1. Because $\epsilon^t > 0$ for each $t \geq 1$, by Algorithm 5 we have $n_s^T \rightarrow \infty$ as $T \rightarrow \infty$. It follows by construction of \mathbf{v}_s^t in Algorithm 5 and by Lemma EC.4 that $\mathbf{v}_s^T \rightarrow \boldsymbol{\pi}_{a(\mathbf{r}^s)}^0$. Step 2. For each $t \geq 1$, construct a solution $\tilde{\mathbf{w}}$ as follows: for each $s \in S$, set $\tilde{w}_{as} = 1$ for $a = a(\mathbf{r}^s)$ and $\tilde{w}_{as} = 0$ for $a \neq a(\mathbf{r}^s)$. Note that $\tilde{\mathbf{w}}$ satisfies constraints (EC.13c)-(EC.13e) for all $t \geq 1$. Next, fix $\tilde{\mathbf{w}}$ as a parameter in PA-T. Then for each t , PA-T simplifies to the following subproblem:

$$\underset{\boldsymbol{\varepsilon}}{\text{minimize}} \quad \sum_{s \in S} |\varepsilon_s| \quad (\text{EC.42a})$$

$$\text{subject to} \quad \left(\sum_{j \in J} v_{s'j}^t r_j^s - c_{a(\mathbf{r}^s)} \right) + \varepsilon_s \geq \left(\sum_{j \in J} v_{s'j}^t r_j^{s'} - c_{a(\mathbf{r}^{s'})} \right), \quad s \in S, s' \in S, \quad (\text{EC.42b})$$

$$\varepsilon_s \geq 0, \quad s \in S. \quad (\text{EC.42c})$$

Let $\tilde{\boldsymbol{\varepsilon}}^t$ be the optimal solution to the above subproblem in round t . It follows from constraint (EC.42b) and (EC.42c) that for each $s \in S$,

$$\tilde{\varepsilon}_s^t = \max_{s' \in S} \left\{ \left(\sum_{j \in J} v_{s'j}^t r_j^s - c_{a(\mathbf{r}^s)} \right) - \left(\sum_{j \in J} v_{s'j}^t r_j^{s'} - c_{a(\mathbf{r}^{s'})} \right), 0 \right\}. \quad (\text{EC.43})$$

Note that $\sum_{j \in J} \pi_{a(\mathbf{r}^{s'}),j}^t r_j^s - c_{a(\mathbf{r}^s)} \leq \sum_{j \in J} \pi_{a(\mathbf{r}^s),j}^t r_j^s - c_{a(\mathbf{r}^s)}$ for all $s \in S$ and $t \geq 1$, by optimality of $a(\mathbf{r}^s)$ with respect to contract \mathbf{r}^s . Because $\mathbf{v}_s^T \rightarrow \boldsymbol{\pi}_{a(\mathbf{r}^s)}^0$ for each $s \in S$ by Step 1, it follows from (EC.43) that $\tilde{\varepsilon}_s^T \rightarrow 0$ for each $s \in S$. Therefore, $Z^T(\tilde{\mathbf{w}}, \tilde{\boldsymbol{\varepsilon}}^T) \rightarrow 0$. By optimality of $(\tilde{\mathbf{w}}^t, \tilde{\boldsymbol{\varepsilon}}^t)$ with respect to PA-T, we have $Z^t(\tilde{\mathbf{w}}^t, \tilde{\boldsymbol{\varepsilon}}^t) \leq Z^t(\tilde{\mathbf{w}}, \tilde{\boldsymbol{\varepsilon}}^t)$ for all $t \geq 1$. Therefore, $Z^T(\tilde{\mathbf{w}}^t, \tilde{\boldsymbol{\varepsilon}}^t) \rightarrow 0$. Step 3. For each $s \in S$ and $t \geq 1$, let $\bar{a}^t(\mathbf{r}^s)$ be the action for which $\bar{\mathbf{w}}_{as}^t = 1$. Suppose by way of contradiction that $\Pr(\bar{a}^T(\mathbf{r}^s) = a(\mathbf{r}^s)) \rightarrow 1$ does not hold for each $s \in S$. Then there exists $\tilde{s} \in S$ and a subsequence $t_b, b = 1, 2, \dots$ such that $\bar{a}^{t_b}(\mathbf{r}^{\tilde{s}}) \neq a(\mathbf{r}^{\tilde{s}})$ for all $b \geq 1$. Next, by constraint (EC.13b), we have

$$\begin{aligned} \varepsilon_{\tilde{s}}^{t_b} &\geq \max_{s' \in S} \left\{ \sum_{j \in J} (v_{s'j}^{t_b} - v_{s'j}^{t_b}) r_j^{\tilde{s}} + c_{\bar{a}(\mathbf{r}^{\tilde{s}})} - c_{\bar{a}(\mathbf{r}^{s'})}, 0 \right\} \\ &\geq \max_{s' \in S} \left\{ \sum_{j \in J} (v_{s'j}^{t_b} - v_{s'j}^{t_b}) r_j^{\tilde{s}} + \underline{c} - \bar{c}, 0 \right\} \end{aligned}$$

for all $b \geq 1$, where the second inequality follows by definition of \underline{c} and \bar{c} . Because $\mathbf{v}_{s'}^{t_b} \rightarrow \boldsymbol{\pi}_{a(\mathbf{r}^{s'})}^0$ as $b \rightarrow \infty$ for all $s' \in S$ by Step 1, and by Assumption EC.3 there exists $s' \in S$ such that $\sum_{j \in J} (\pi_{a(\mathbf{r}^{s'}),j}^0 - \pi_{a(\mathbf{r}^{\tilde{s}}),j}^0) r_j^{\tilde{s}} + \underline{c} - \bar{c} > 0$, it follows that $\lim_{b \rightarrow \infty} \varepsilon_{\tilde{s}}^{t_b} > 0$. Hence $\lim_{b \rightarrow \infty} Z^{t_b}(\tilde{\mathbf{w}}^{t_b}, \tilde{\boldsymbol{\varepsilon}}^{t_b}) > 0$, which implies $\lim_{T \rightarrow \infty} Z^T(\tilde{\mathbf{w}}^t, \tilde{\boldsymbol{\varepsilon}}^t) > 0$. However, by Step 2, $Z^T(\tilde{\mathbf{w}}^t, \tilde{\boldsymbol{\varepsilon}}^t) \rightarrow 0$, which yields a contradiction. Therefore, $\Pr(\bar{a}^T(\mathbf{r}^s) = a(\mathbf{r}^s)) \rightarrow 1$ for all $s \in S$. By definition of $\bar{a}^t(\mathbf{r}^s)$, this implies $\Pr(\bar{w}_{as}^T = 1) \rightarrow 1$ for $a = a(\mathbf{r}^s)$ and $\Pr(\bar{w}_{as}^T = 0) \rightarrow 1$ for $a \neq a(\mathbf{r}^s)$, as desired. Step 4. We now show $\hat{\boldsymbol{\pi}}^T \rightarrow \boldsymbol{\pi}^0$. For convenience, let $S_a = \{s | a(\mathbf{r}^s) = a\}$. Then by construction of $\hat{\boldsymbol{\pi}}^t$ from Algorithm 5,

$$\hat{\pi}_{aj}^t = \frac{\sum_{s \in S} v_{sj}^t n_s^t \bar{w}_{as}^t}{\sum_{s \in S} n_s^t \bar{w}_{as}^t} = \frac{\sum_{s \in S_a} v_{sj}^t n_s^t \bar{w}_{as}^t + \sum_{s \in S \setminus S_a} v_{sj}^t n_s^t \bar{w}_{as}^t}{\sum_{s \in S_a} n_s^t \bar{w}_{as}^t + \sum_{s \in S \setminus S_a} n_s^t \bar{w}_{as}^t},$$

for all (a, j) . By Step 3, $\Pr(\bar{w}_{as}^T = 1) \rightarrow 1$ for $s \in S_a$ and $\Pr(\bar{w}_{as}^T = 0) \rightarrow 1$ for $s \in S \setminus S_a$. Further, by Step 1, $\mathbf{v}_s^T \rightarrow \boldsymbol{\pi}_{a(r^s)}^0$. It follows that for each $j \in J$, $v_{s_j}^T \bar{w}_{as}^T \rightarrow \pi_{aj}^0$ for $s \in S_a$, and $v_{s_j}^T \bar{w}_{as}^T \rightarrow 0$ for $s \in S \setminus S_a$. Therefore,

$$\text{plim}_{T \rightarrow \infty} \hat{\boldsymbol{\pi}}_{aj}^T = \text{plim}_{T \rightarrow \infty} \frac{\sum_{s \in S_a} v_{s_j}^T n_s^T \bar{w}_{as}^T + \sum_{s \in S \setminus S_a} v_{s_j}^T n_s^T \bar{w}_{as}^T}{\sum_{s \in S_a} n_s^T \bar{w}_{as}^T + \sum_{s \in S \setminus S_a} n_s^T \bar{w}_{as}^T} = \text{plim}_{T \rightarrow \infty} \frac{\pi_{aj}^0 \sum_{s \in S_a} n_s^T}{\sum_{s \in S_a} n_s^T} = \pi_{aj}^0.$$

for all (a, j) . Because $\hat{\boldsymbol{\pi}}_{aj}^T \rightarrow \pi_{aj}^0$ for all (a, j) , it follows that $\hat{\boldsymbol{\pi}}^T \rightarrow \boldsymbol{\pi}^0$. \square

Proof of Corollary EC.2. It suffices to show $|U(\hat{\mathbf{r}}^T) - U(\mathbf{r}^*)| \rightarrow 0$. The proof proceeds in two steps. First, we show $\hat{U}(\mathbf{r}) \rightarrow U(\mathbf{r})$ for any $\mathbf{r} \in R$. Second, we prove the main result. Step 1. Note

$$\begin{aligned} |U(\hat{\mathbf{r}}) - U(\mathbf{r}^*)| &= |U(\hat{\mathbf{r}}) - \hat{U}(\hat{\mathbf{r}}) + \hat{U}(\hat{\mathbf{r}}) - \hat{U}(\mathbf{r}^*) + \hat{U}(\mathbf{r}^*) - U(\mathbf{r}^*)| \\ &\leq |U(\hat{\mathbf{r}}) - \hat{U}(\hat{\mathbf{r}})| + |\hat{U}(\hat{\mathbf{r}}) - \hat{U}(\mathbf{r}^*)| + |\hat{U}(\mathbf{r}^*) - U(\mathbf{r}^*)| \end{aligned}$$

where the equality follows from adding and subtracting $\hat{U}(\hat{\mathbf{r}})$ and $\hat{U}(\mathbf{r}^*)$, and the inequality follows from the triangle inequality. Next, note that by the definitions of $\hat{U}(\mathbf{r})$ (given above) and $U(\mathbf{r})$ (given in (EC.11)), $\hat{\boldsymbol{\pi}}^T \rightarrow \boldsymbol{\pi}^0$ implies $\hat{U}(\mathbf{r}) \rightarrow U(\mathbf{r})$ for any $\mathbf{r} \in R$. Step 2. By Step 1, $|U(\hat{\mathbf{r}}^T) - \hat{U}(\hat{\mathbf{r}}^T)| \rightarrow 0$ and $|\hat{U}(\mathbf{r}^*) - U(\mathbf{r}^*)| \rightarrow 0$. It remains to show that $|\hat{U}(\hat{\mathbf{r}}^T) - \hat{U}(\mathbf{r}^*)| \rightarrow 0$. Suppose not. Then there exists a subsequence $t_b, b \geq 1$ such that $|\hat{U}(\hat{\mathbf{r}}^{t_b}) - \hat{U}(\mathbf{r}^*)| > 0$ for all $b \geq 1$. Further, because for all $t \geq 1$, $\hat{\mathbf{r}}^t \in \text{argmax}_{\mathbf{r} \in R} \hat{U}(\mathbf{r})$ by definition, we have $\hat{U}(\hat{\mathbf{r}}^{t_b}) > \hat{U}(\mathbf{r}^*)$ for all $b \geq 1$. Letting $b \rightarrow \infty$ and noting $\hat{U}(\mathbf{r}) \rightarrow U(\mathbf{r})$ for all $\mathbf{r} \in R$ yields $\text{plim}_{b \rightarrow \infty} \hat{U}(\hat{\mathbf{r}}^{t_b}) > U(\mathbf{r}^*)$. It is straightforward to obtain a contradiction to the preceding inequality using $\hat{U}(\mathbf{r}) \rightarrow U(\mathbf{r})$ from Step 1, and $\mathbf{r}^* = \text{argmax}_{\mathbf{r} \in R} U(\mathbf{r})$ by definition of \mathbf{r}^* . \square