

Counterfactual Data-Fusion for Online Reinforcement Learners

Andrew Forney

University of California, Los Angeles
580 Portola Plaza
Los Angeles, CA 90095
forns@cs.ucla.edu

Judea Pearl

University of California, Los Angeles
580 Portola Plaza
Los Angeles, CA 90095
judea@cs.ucla.edu

Elias Bareinboim

Purdue University
305 N. University Street
West Lafayette, IN 47907
eb@purdue.edu

ABSTRACT

Online learning agents possess the capacity to learn from both personalized experimentation as well as the observed behaviors of other agents interacting with the environment. However, data collected through these different modalities may not be naively combined due to changes in the decision-making context, including factors that may be unobserved. The data-fusion problem addresses how information collected under such disparate conditions (observationally, experimentally, and counterfactually) can be combined to yield more informative results than the independent datasets alone. The present work provides a recipe for combining multiple datasets to accelerate learning in a variant of the Multi-Armed Bandit problem with Unobserved Confounders (MABUC). We demonstrate this data-fusion approach through an enhanced Thompson Sampling bandit player, and support its efficacy with extensive simulations.

CCS Concepts

•Computing methodologies → Online learning settings;

Keywords

Learning Agent Capabilities; Data-Fusion; Counterfactual Logic

1. INTRODUCTION

As active learning agents become increasingly integrated into real-world environments, they gain new sources of information related to their tasks at hand. Not only do these agents possess the ability to interact with their environments (choosing actions, receiving feedback on the quality of their choices, and then modifying future actions accordingly), they may also observe other agents doing the same. However, with opportunities to adjust policies from sources other than personal experimentation come new challenges of “transfer” in learning. In particular, agents should be wary of how observed behavior generalizes (i.e., transfers) to them, how these observations should be combined with the agent’s own experience, and how such a combination can be robustly

maintained in the face of changing environmental factors.

In this work, we consider how data collected by an online agent under various conditions (e.g., experimental vs. non-experimental settings) can be combined to improve performance in a reinforcement learning task. This challenge is not without precedent, as recent studies have investigated dataset transportability, though in offline domains [3]. Others have studied scenarios in which agents learn from expert teachers in the inverse reinforcement learning problems [1, 9]. Furthermore, recent results from causal analysis have explored personalized decision-making for agents in reinforcement learning tasks [2, 18], but did not provide a complete framework for fusing the disparate dataset classes that we will introduce shortly.

Environments for which an agent (1) observes all state variables and (2) possesses a fully specified model (in which all factors relating contexts, actions, and their associated rewards are known) are trivial from a learning transfer perspective; in such scenarios, collected data is homogeneous because all factors that may introduce bias between samples can be controlled. Conversely, in this paper, we focus on the challenges that arise due to *unobserved confounders* (UCs), namely, unmeasured variables that influence an agent’s natural action choice as well as the feedback from that action¹. Such factors are particularly subtle when left uncontrolled due to their invisible nature and potential to introduce *confounding bias* [12, Chs. 3,6].

Because our agent’s goal is to quickly learn an optimal policy by consolidating data collected from observing other agents and data collected through its own experience, UCs pose a fundamental challenge: the results from *seeing* another agent performing an action are not necessarily *interchangeable* with those from *doing* the action itself. As such, throughout this paper, we will differentiate three classes of data that may be employed by an autonomous agent to inform its decision-making:

1. **Observational data** is gathered through passive examination of the actions and rewards of agents other than the actor, but for whom the actor is assumed to be exchangeable.
2. **Experimental data** is gathered through randomization, or from fixed policies that are not reactive to the environmental state.

Appears in: *Proceedings of the 1st Workshop on Transfer in Reinforcement Learning (TiRL) at the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, A. Costa, D. Precup, M. Veloso, M. Taylor (chairs), May 8–9, 2017, São Paulo, Brazil.

¹The natural (or “observational”) regime represents when the agent is not influenced by external events and makes decisions following her standard protocol [12, Ch. 3].




Figure 1: Plots of MAB algorithms performance vs. RDC in the Greedy Casino scenario. Note that all algorithms but TS^{RDC} experience linear regret.

3. **Counterfactual data** represents the rewards associated with actions under a particular (or “personalized”) configuration of the UCs.

In the remainder of this work, we demonstrate how these data types can be fused to facilitate learning in a variant of the Multi-Armed Bandit problem with Unobserved Confounders (MABUC), first discussed in [2, 11, 5]. In traditional bandit instances (e.g., [14, 10, 6, 15, 4]), an agent is faced with $K \in \mathbb{N}$, $K \geq 2$ discrete action choices (often called “arms”), each with its own, independent, and initially unknown reward distribution. The agent’s task is to maximize cumulative rewards over a series of rounds, which requires learning about the underlying reward distributions associated with each arm. In the MABUC (formalized shortly), agents are faced with the same task, except that UCs modify the agent’s arm-choice predilections and payout rates at each round, and the dimensionality and functional form of the UCs are unknown.

Though the data-fusion problem is an ongoing exploration in the data sciences [3], this paper is the first to study online learning techniques in MABUC settings that combine data sampled under disparate conditions. Specifically, our contributions are as follows:

1. Using counterfactual calculus, we prove that counterfactual quantities can be empirically estimated by active agents (Sec. 4).
2. We demonstrate how observational, experimental, and counterfactual datasets can be combined through a heuristic for MABUC agents (Sec. 4).
3. We then develop a variant of the Thompson Sampling algorithm that implements this new heuristic, and run extensive simulations demonstrating its faster convergence rates compared to the current state-of-the-art (Sec. 5).

2. THE GREEDY CASINO REVISITED

In this section, we consider an expanded version of the Greedy Casino problem introduced in [2]. In its new floor’s configuration, the Greedy Casino is considering four new themed slot-machines (instead of the two used in the previous version) and wishes to make them as lucrative as possible. After running a battery of preliminary tests, the

(a)	$D = 0$		$D = 1$	
	$B = 0$	$B = 1$	$B = 0$	$B = 1$
$E[y_1 X, B, D]$				
$X = 0$	*0.20	0.30	0.50	0.60
$X = 1$	0.60	*0.20	0.30	0.50
$X = 2$	0.50	0.60	*0.20	0.30
$X = 3$	0.30	0.50	0.60	*0.20

(b)	$E[y_1 X]$	$E[y_1 do(X)]$
	$X = 0$	0.20
$X = 1$	0.20	0.40
$X = 2$	0.20	0.40
$X = 3$	0.20	0.40

Table 1: (a) Payout rates decided by reactive slot machines as a function of arm choice X , sobriety D , and machine conspicuousness B . Players’ natural arm choices ($f_x = B + 2D$) under D, B are indicated by asterisks. (b) Payout rates according to the observational, $E[y_1|X]$, and experimental $E[y_1|do(X)]$, distributions, where $Y = y_1$ represents winning (shown in the table).

casino executives discover that two traits in particular predict which of the four machines that a gambler is likely to play: whether or not the machines are blinking (denoted $B \in \{0, 1\}$), and whether or not the gambler is drunk (denoted $D \in \{0, 1\}$). After consulting with a team of psychologists and statisticians, the casino learns that any arbitrary gambler’s natural machine choice can be modeled by the structural equation [12]: $X \leftarrow f_X(B, D) = B + 2 * D$ if the four machines are indexed as $X \in \{0, 1, 2, 3\}$. The casino also learns that its patrons have an equal chance of being drunk or not (i.e., $P(D = 1) = 0.5$) and decide to program their new machines to blink half of the time (i.e., $P(B = 1) = 0.5$).

To prevent casinos from exploiting their patrons, a new gambling law stipulates that all slot machines in the state must maintain a minimum 30% win rate. Wishing to leverage their new discovery about gamblers’ machine choice predilections while conscious of this law, the casino implements a reactive payout strategy for their machines, which are equipped with sensors to determine if their gambler is drunk or not (assume that the sensors are perfect at making this determination). As such, the machines are programmed with the payout distribution illustrated in Table 1a.

After the launch of the new slot machines, some observant gamblers note that players appear to be winning only 20% of the time, and report their suspicions to the state gambling commission. An investigator is then sent to the casino to determine the merit of these complaints, and begins recruiting random gamblers from the casino floor to play at randomly selected machines, despite the players’ natural predilections. Surprisingly, he finds that players in this experiment win 40% of the time, and declares that the casino has committed no crime. Meanwhile, the casino continues to exploit players’ gambling predilections, paying them 10% less than the law-mandated minimum. Plainly, gamblers are unaware of being manipulated by the UCs B, D , and of the predatory payout policy that the casino has constructed around them. The collected data is summarized in Table 1b; the second column ($E[y_1|X]$) represents the observations drawn from the casino’s floor while the third ($E[y_1|do(X)]$) represents the randomized experiment performed by the state

investigator (both with large sample sizes).

In an attempt to find a better gambling strategy, an observant habitué decides to run a battery of experiments using standard MAB algorithms (e.g., ϵ -greedy, UCB, Thomson Sampling) as well as an algorithm following an approach presented in [2] known as the Regret Decision Criterion (RDC) (reviewed in the next section). Importantly, the RDC agent lacks the capability to identify and observe the UCs. The results of her experiments are depicted in Fig. 1. She notes, somewhat surprised, that all algorithms which ignore the influence of the UCs (i.e., all but RDC) perform equivalently to the randomized experiment conducted by the investigator. Noting the differences in the payout rates between the observational and experimental data, she ponders how this can be the case and how she might use these datasets to improve her gambling strategy and winnings.

3. BACKGROUND

In this section, we formalize the MABUC problem in the language of Structural Causal Models (SCMs), which will allow us to articulate the notions of observational, experimental, and counterfactual distributions as well as formalize the problem of confounding due to the influence of UCs.

Each SCM M is associated with a causal diagram G and encodes a set of endogenous (or observed) variables V and exogenous (or unobserved) variables U ; edges in G correspond to functional relationships relative to each endogenous variable $V_i \in V$, namely, $V_i \leftarrow f_i(PA_i, U_i)$, where $PA_i \subseteq V \setminus V_i$ and $U_i \subseteq U$; and a probability distribution over the exogenous variables $P(U = u)$.

Each M induces a set of: (1) observational distributions $P(X = x)$ for $X \subseteq V$, which represent the “natural” world, without external interventions; (2) experimental (a.k.a. interventional) distributions $P(Y = y|do(X = x))$ for $X, Y \subseteq V$, which represent the world in which X is forced to the value x despite any causal influences that would otherwise functionally determine its value in the natural setting; and (3) counterfactual distributions, defined next [12, pp. 203-205].²

Definition 3.1. (Counterfactual) [12, pp. 204] Let X and Y be two subsets of endogenous variables in V . The counterfactual sentence “ Y would be y (in situation $U = u$), had X been x ” is interpreted as the equality with $Y_x(u) = y$, where $Y_x(u)$ encodes the solution for Y in a structural system where for every $V_i \in X$, the equation f_i is replaced with the constant x .

For example, the counterfactual expectation $E[Y_x = y|X = x']$ is well-defined, even when $x \neq x'$, and is read “The expectation that $Y = y$ had X been x given that X was observed to be x' ”. Despite being logically valid statements, counterfactual quantities must be estimated from either a fully specified model, or, in the absence of such, from data. In offline settings, however, counterfactual quantities are not empirically estimable (namely, when the antecedent of the counterfactual contradicts the observed value), except in some special cases [12, Ch. 9]. The reason is that if we submitted the subject to condition $X = x$, we cannot go back in time before exposure to treatment and submit the same subject to a new condition $X = x'$. As is well understood in the

²For a comprehensive review of SCMs, we refer readers to [12, Ch. 7].

causal inference literature, this procedure is not the same as first exposing a random unit to condition $X = x'$ since the ones who initially were inclined to act as $X = x$ are somehow different than the randomly selected subject. That said, we will prove (in Sec. 4) that online learning agents possess the means to estimate counterfactuals directly.

In practice, the observational and experimental distributions can be estimated through procedures known as *random sampling* and *random experimentation*, respectively. Confounding bias emerges when UCs are present and can be seen through the difference between these two distributions, $P(Y|do(X = x)) - P(Y|X = x)$ [12, Ch. 6]. The absence of UCs implies that $P(Y|do(X = x)) = P(Y|X = x)$, which allows random sampling (instead of a randomized experiment) to estimate the experimental distribution.

The contrast between observational and experimental data is mirrored in the distinction between *actions* (which represent reactive “choices” resulting from an agents’ environments, beliefs, and other causes) and *acts* (which represent deliberate choices resulting from rational decision-making or interventions that sever the causal influences of the system [12, Sec. 4.1]). To tie these concepts to the MABUC problem, one important tool introduced in [2] is what is known as the agent’s *intent*.

Definition 3.2. (Intent) Consider a SCM M and an endogenous variable $X \in V$ that is amenable to external interventions and is (naturally) determined by the structural function $f_x(PA_x, U_x)$, where $PA_x \subseteq V$ represents the observed parents of X , and $U_x \subseteq U$ are the UCs of X . After realization $PA_x = pa_x$ and $U_x = u_x$ (without any external intervention), the realization $I = f_x(pa_x, u_x)$ is called the agent’s intent.

Thus, we can consider intent as an agent’s chosen *action* before its execution, serving as a proxy for any influencing UCs.³ To ground these notions, consider again the Greedy Casino example in which the gamblers’ intents are enacted on the unperturbed casino floor, but are then averaged over during the investigator’s randomized study.

We can now put these observations together and explicitly define the MABUC problem:

Definition 3.3. (K-Armed Bandits with Unobserved Confounders) A K -Armed bandit problem ($K \in \mathbb{N}, K \geq 2$) with unobserved confounders (MABUC, for short) is defined as a model M with a reward distribution over $P(u)$ where, for each round $0 < t < T, t \in \mathbb{N}$:

1. *Unobserved confounders:* U_t represents the unobserved variable encoding the payout rate and unobserved influences to the propensity to choose arm x_t at round t .
2. *Intent:* $I_t \in \{i_1, \dots, i_k\}$ represents the agent’s intended arm choice at round t (prior to its final choice, X_t) such that $I_t = f_i(pa_{x_t}, u_t)$.
3. *Policy:* $\pi_t \in \{x_1, \dots, x_k\}$ denotes the agent’s decision algorithm as a function of its history (discussed shortly)

³Def. 3.2 does not require that all of the agent’s influencing factors be measured or acknowledged. This definition accommodates the fact that an agent’s decisions can be influenced by unknown factors, an observation that is not new to the cognitive sciences [16].




Figure 2: Model for each round of the MABUC decision process. Solid nodes denote observed variables and open nodes represent unobserved variables. The square node indicates a decision point made by the agent’s strategy.

and current intent, $f_\pi(h_t, i_t)$ ⁴.

4. *Choice:* $X_t \in \{x_1, \dots, x_k\}$ denotes the agent’s final arm choice that is “pulled” at round t , $x_t = f_x(\pi_t)$.
5. *Reward:* $Y_t \in \{0, 1\}$ represents the Bernoulli reward (0 for losing, 1 for winning) from choosing arm x_t under UC state u_t as decided by $y_t = f_y(x_t, u_t)$.

The graphical model in Fig. 2 represents a prototypical MABUC (Def. 3.3). We also add a graphical representation of the agent’s history H_t , a data structure containing the agent’s observations, experiments, and counterfactual experiences up to time step t . The means by which these datasets can be used in the agent’s decision-making are explored at length in the next section. In summary, at every round t of MABUC, the unobserved state u_t is drawn from $P(u)$, which then decides i_t , which is then considered by the strategy π_t in concert with the game’s history h_t ; the strategy makes a final arm choice, which is then pulled, as represented by x_t , and the reward y_t is revealed.

Based on this definition, the regret decision criterion was developed.

Definition 3.4. (Regret Decision Criterion (RDC)) [2] In a MABUC instance with arm choice X , intent $I = i$, and reward Y , agents should choose the action a that maximizes their intent-specific reward, or formally:

$$\operatorname{argmax}_a E[Y_{X=a} = 1 | X = i] \quad (1)$$

In brief, RDC prescribes that the arm $X = a$ that maximizes the expected value of reward Y having conditioned on the intended arm $X = i$ should be selected, even when $a \neq i$.

4. FUSING DATASETS

Suppose our agent assumes the role of a gambler in the Greedy Casino (Sec. 2) and possesses (1) observations of

⁴This representation clearly distinguishes observational and experimental settings – π_t copies i_t in the former, but ignores i_t and listens instead to a random device (e.g., a coin toss) in the latter.

arm choices and payouts from other players in the casino, (2) the randomized experimental results from the state investigator, and (3) the knowledge to use intent in its decision-making for choosing arms by the Regret Decision Criterion (RDC). In other words, the agent begins the MABUC problem with large samples of observations ($E[Y|X]$) and experimental results ($E[Y|do(X)]$), and will maximize the counterfactual RDC ($E[Y_{X=a} = 1 | X = i]$) because it recognizes the presence of UCs (viz. $E[Y|X] \neq E[Y|do(X)]$; see Table 1b). We note that the observational and experimental data available to our agent contains information about its environment, but cannot simply be incorporated into the counterfactual maximization criteria (viz. $E[Y|X] \neq E[Y|do(X)] \neq E[Y_x|x]$; see Table 1). So, the agent can choose to either discard its observations and experiments, and simply gamble by the tenets of RDC, or combine them in an informative way. This section explores the latter option.

Relating the Datasets

Note that we can write the experimental quantity $E[Y|do(X = x)]$ in the counterfactual notation $E[Y_{X=x}] = E[Y_x]$ reading “The expected value of Y had X been x .” Relating this quantity to the RDC, by the law of total probability we can write $E[Y_x]$ as a weighted average of the reward associated with arm x across all intent conditions, namely:

$$E[Y_x] = E[Y_x|x_1]P(x_1) + \dots + E[Y_x|x_K]P(x_K) \quad (2)$$

Examining Eq. 2, we see that the equation is composed of expressions from our agent’s three datasets (observational, experimental, and counterfactual). By definition, the LHS of the equation ($E[Y_x]$) is drawn from the experimental dataset. On the RHS, we have two types of quantities. Expressions of the form $E[Y_x|x]$ for which $x = x'$ are observational by the *consistency axiom* [12, pp. 229], because, despite being written in counterfactual notation, the observation (x') and antecedent (x) are the same action; in other words, the hypothesized action and observed action are the same, and so $E[Y_x|x] = E[Y|x]$. However, expressions of the form $E[Y_x|x']$ for which $x \neq x'$ are indeed counterfactual.

In general, evaluating counterfactuals empirically is not possible, except for some special cases, such as when the action X is binary [12, Ch. 9]. However, RDC asserts that if one preempts the agent’s decision process when the intent $I = i$ is about to become a decision (X), i still encodes information about the UCs U_i (because $i = f_i(PA_X)$). This implies that randomizing within intent conditions can lead to the computation of the counterfactual given by RDC, which is a special counterfactual also called *the effect of treatment on the treated* (ETT) [12, Ch. 7].

In order for us to exploit the properties of this equivalence to improve the performance of RDC agents in the MABUC setting, we first demonstrate that RDC indeed measures the counterfactual quantity of the ETT.

Theorem 4.1. *The counterfactual ETT is empirically estimable for arbitrary action-choice dimension (i.e., $|X| = k$ for $k \geq 2$) when agents condition on their intent $I = i$ and estimate the response Y to their final action choice $X = a$.*

PROOF. We start by writing the corresponding ETT ex-

pansion and note that,

$$E[Y_{X=a}|X=i] \quad (3)$$

$$= \sum_{i'} E[Y_{X=a}|X=i, I=i']P(I=i'|X=i) \quad (4)$$

$$= \sum_{i'} E[Y_{X=a}|I=i']P(I=i'|X=i) \quad (5)$$

$$= \sum_{i'} E[Y|do(X=a), I=i']P(I=i'|X=i) \quad (6)$$

$$= \sum_{i'} E[Y|do(X=a), I=i']1(i'=i) \quad (7)$$

$$= E[Y|do(X=a), I=i] \quad (8)$$

Eq. (4) expands the ETT using the law of total probability to sum over all intent conditions. Eq. (5) follows from the conditional independence $Y_x \perp\!\!\!\perp X|I$ that holds, allowing us to remove $X=i$. Eq. (6) follows because intent, by definition, satisfies the counterfactual back-door criterion [13, Theorem 4.3.1]. Eqs. (7,8) follow from the fact that, observationally, an agent’s final arm choice will always coincide with their intent (i.e., $P(i|x) = 1 \forall i=x, 0$ otherwise), which nullifies all summed expressions where the two differ. \square

Because Eq. (8) is an interventional quantity, we have shown that the ETT, a counterfactual expression, can be estimated empirically through intent-specific randomization.⁵ The main advantages of this, now proven, equivalence are threefold: (1) the empirical estimation of previously unidentifiable counterfactual quantities presents opportunities for further exploration in causal analysis, (2) the ETT’s prescription for integrating experimental and observational data (see Eq. 2) permits a now interventional data-fusion strategy when such data is available, and (3) data points sampled by the agent using intent-specific decision-making are counterfactual in nature, and should therefore be added to the agent’s counterfactual history. As such, we now consider that procedures implementing RDC-type randomization can record intent-specific arm rewards in a data structure similar to Fig. 3. A second consequence of recording arm-intent-specific payouts in this fashion is that observational data may be substituted directly into cells for which the final arm choice and intent agree (see reference to consistency axiom below Eq. 2).

Strategies for Online Agents

Having illustrated how the different datasets relate to our agent in the MABUC setting, we now consider that the counterfactual expressions in Eq. 2 must be learned by our agent and are not known at the start of the game. Because of this finite-sample concern, we propose several learning strategies that exploit the datasets’ relationship while managing the uncertainty implicit in a MAB learning scenario.

Strategy 1: Cross-Intent Learning. Consider Eq. 2 once more. This holds for every arm $X=x$, which induces a system of equations as shown in Fig. 3. Consider a single cell in this system, say $E[Y_{x_r}|x_w]$, which we can solve and

⁵It is understood that the ETT can be computed for binary decisions or when the backdoor criterion holds [12], but it was not believed to be estimable for arbitrary dimensions prior to RDC-randomization.

	$I=x_1$	$I=x_2$...	$I=x_K$
$X=x_1$	$E[Y_{x_1} x_1]$	$E[Y_{x_1} x_2]$...	$E[Y_{x_1} x_K]$
$X=x_2$	$E[Y_{x_2} x_1]$	$E[Y_{x_2} x_2]$...	$E[Y_{x_2} x_K]$
\vdots	\vdots	\vdots	\ddots	...
$X=x_K$	$E[Y_{x_K} x_1]$	$E[Y_{x_K} x_2]$	\vdots	$E[Y_{x_K} x_K]$

Figure 3: Counterfactual reward-history table as a cross of arm choice and intent.

rewrite as:

$$E_{XInt}[Y_{x_r}|x_w] = [E[Y_{x_r}] - \sum_{i \neq w}^K E[Y_{x_r}|x_i]P(x_i)]/P(x_w) \quad (9)$$

This strategy provides a systematic way of learning about arm payouts across intent conditions, which is desirable because an arm pulled under one intent condition provides knowledge about the payouts of that arm under other intent conditions. This can be depicted graphically, as shown by row B in Fig. 3 – information about Y_{x_r} flows from intent conditions $x_i \neq x_w$ to intent x_w .

Strategy 2: Cross-Arm Learning. Consider any three arms, x_r, x_s, x_w such that $r \neq \{s, w\}$ and assume we are interested in estimating the value of $E[Y_{x_r}|x_w]$ (our query, for short). Considering again the equations induced by Eq. (2), we have,

$$E[Y_{x_r}] = \sum_i^K E[Y_{x_r}|x_i]P(x_i) \quad (10)$$

$$E[Y_{x_s}] = \sum_i^K E[Y_{x_s}|x_i]P(x_i) \quad (11)$$

Note that each of Eqs. (10, 11) share the same intent priors on our query intent $P(x_w)$, so we can solve for $P(x_w)$ in both equations using simple algebra, which yields,

$$P(x_w) = \frac{E[Y_{x_r}] - \sum_{i \neq w}^K E[Y_{x_r}|x_i]P(x_i)}{E[Y_{x_r}|x_w]} \quad (12)$$

$$= \frac{E[Y_{x_s}] - \sum_{i \neq w}^K E[Y_{x_s}|x_i]P(x_i)}{E[Y_{x_s}|x_w]} \quad (13)$$

Using Eq. (12) and solving for the query in terms of our paired arm x_s , $\forall r \neq s$ we have

$$E[Y_{x_r}|x_w] = \frac{[E[Y_{x_r}] - \sum_{i \neq w}^K E[Y_{x_r}|x_i]P(x_i)]E[Y_{x_s}|x_w]}{E[Y_{x_s}] - \sum_{i \neq w}^K E[Y_{x_s}|x_i]P(x_i)} \quad (14)$$

Eq. (14) illustrates that any non-diagonal cell from the table in Fig. 3 can be estimated through pairwise arm comparisons with the same intent. Put differently, Eq. (14) allows our agent to estimate $E[Y_{x_r}|x_w]$ from samples in which any arm $x_s \neq x_r$ was pulled under the same intent x_w .

In practice, the online nature of the problem can make some of these pairwise computations noisy due to sampling

variability when x_r is an infrequently explored arm. To obtain a more robust estimate of the target quantity, this pairwise comparison can be repeated between the query arm and all other arms with the same intent, and then pooled together. (Graphically, this can be seen as information about $Y_{x_r}|x_w$ flowing from arm $x_s \neq x_r$ to x_r (under intent x_w) – column C in Fig. 3.)

One such pooling strategy is to take the *inverse-variance-weighted* average.⁶ Formally, we can consider a function $E[Y_{x_r}|x_w] = h_{XArm}(x_r, x_w, x_s)$ such that h_{XArm} performs the empirical evaluation of the RHS of Eq. (14). Additionally, let $\sigma_{x_i}^2$ indicate the empirical payout variance for each arm-intent condition (as from the reward successes and failures captured by the agent in Table 3). To estimate our query from all other arms in the same intent through inverse-variance weighting, we have our now complete, second heuristic:

$$E_{XArm}[Y_{x_r}|x_w] = \frac{\sum_{i \neq r}^K h_{XArm}(x_r, x_w, x_i) / \sigma_{x_i, x_w}^2}{\sum_{i \neq r}^K 1 / \sigma_{x_i, x_w}^2} \quad (15)$$

Strategy 3: The Combined Approach. The payout estimates for a MABUC algorithm using RDC (Fig. 3) can be estimated from three different sources:

1. $E_{samp}[Y_{x_r}|x_w]$, the *sample* estimates collected by the agent during the execution of the algorithm.
2. $E_{XInt}[Y_{x_r}|x_w]$, the computed estimate using *cross-intent* learning.
3. $E_{XArm}[Y_{x_r}|x_w]$, the computed estimate using *cross-arm* learning.

Naturally, these three quantities can be combined to obtain a more robust and stable estimate to the target query.

We employ an inverse-variance weighting scheme so as to leverage these three estimators, and so we must formulate a metric for the payout variance associated with each strategy’s computed estimate. To do so, we define an average variance for each strategy, which is the average over each sample estimate’s variance (i.e., $\sigma_{x_i}^2$) used in the computation. Specifically, for the cross-arm approach (Eq. 15), we have two summations over sample payout estimates $E[Y_{x_r}|x_i]$, $E[Y_{x_s}|x_i] \forall i \neq w$ which involve $2(K-1)$ terms, plus the numerator’s $E[Y_{x_s}|x_w]$, giving us a total of $2(K-1) + 1 = 2K-1$ variances to average. The same is true for the cross-intent approach (Eq. 9), which involves $K-1$ sample variances to average. When estimating $E[Y_{x_r}|x_w]$, we can write the corresponding variances:

$$\sigma_{XArm}^2 = \frac{1}{2K-1} \left[\sum_{i \neq w}^K \sigma_{x_r, x_i}^2 + \left[\sum_{i \neq w}^K \sigma_{x_s, x_i}^2 + \sigma_{x_s, x_w}^2 \right] \right]$$

$$\sigma_{XInt}^2 = \frac{1}{K-1} \sum_{i \neq w}^K \sigma_{x_r, x_i}^2$$

Finally, to estimate $E[Y_{x_r}|x_w]$ using our combined ap-

⁶This strategy follows from the fact that we have Bernoulli rewards for each arm-intent condition, and as the number of samples increases for these distributions, the variance diminishes, meaning that arm-intent conditions with smaller variances are more reliable than those with larger ones.

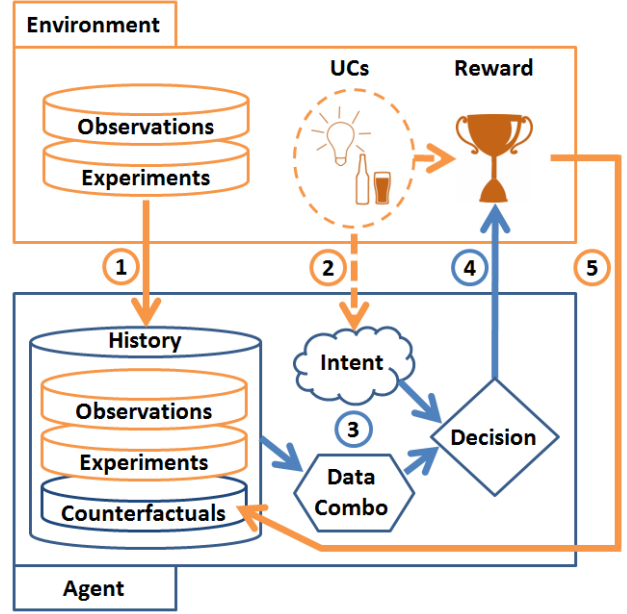


Figure 4: Illustrated data-fusion process.

proach, we have:

$$\alpha = E_{samp}[Y_{x_r}|x_w] / \sigma_{x_r, x_w}^2 + E_{XInt}[Y_{x_r}|x_w] / \sigma_{XInt}^2 + E_{XArm}[Y_{x_r}|x_w] / \sigma_{XArm}^2$$

$$\beta = 1 / \sigma_{x_r, x_w}^2 + 1 / \sigma_{XInt}^2 + 1 / \sigma_{XArm}^2$$

$$E_{combo}[Y_{x_r}|x_w] = \frac{\alpha}{\beta} \quad (16)$$

To visualize this data-fusion process in the context of a reinforcement learning problem, consider the diagram in Figure 4. (1) In this scenario, we consider that our agent has collected large samples of experimental and observational data from its environment (e.g., in the Greedy Casino, the agent might observe other gamblers to comprise its observational data and incorporate experimental findings from the state investigator’s report). (2) Unobserved confounders are realized in the environment, though their labels and values are unknown to the agent. (3) From these UCs and any other observed features in the environment, the agent’s heuristics suggest an action to take, i.e., its intent. With its intent known, the agent combines the data in its history (in this work, by the prescription of Strategy 3 above) to better inform its decision-making. (4) Based on its intent and combined history, the agent commits to a final action choice. (5) The action’s response in the environment (i.e., its reward) is observed, and the collected data point is added to the agent’s counterfactual history (as a consequence of Theorem 4.1).

5. SIMULATIONS & RESULTS

In this section, we validate the efficacy of the strategies discussed previously through simulations. To make a fair comparison to previous MABUC bandit players, we will follow the first implementation of RDC that used Thompson Sampling (TS) as its basis, embedding the strategies described in the previous section within a TS player called

(TS^{RDC+}). We note that after moving from traditional to counterfactual-based decision-making we moved from optimal arm-choice nonconvergence to convergence in the MABUC setting (e.g., Fig. 1); now, the goal is to accelerate convergence.

In brief, TS^{RDC+} agents perform the following at each round:

1. Observe the intent i_t from the current round’s realization of UCs, u_t .
2. Sample $\hat{E}_{samp}[Y_{x_r}|i_t]$ from each arm’s (x_r) corresponding intent-specific beta distribution $\beta(s_{x_r,i_t}, f_{x_r,i_t})$ ⁷ in which s_{x_r,i_t} is the number of successes (wins) and f_{x_r,i_t} is the number of failures (losses).
3. Compute each arm’s i_t -specific score using the combined datasets via Strategy 3 (Eq. 16).
4. Choose the arm, x_a , with the highest score computed in previous step.
5. Observe result (win / loss) and update $\hat{E}_{samp}[Y_{x_a}|i_t]$.

Procedure. Simulations were performed on the 4-arm MABUC problem, with results averaged across $N = 1000$ Monte Carlo repetitions, each $T = 3000$ rounds in duration. To illustrate the robustness of each proposed strategy, we performed simulations spanning across a wide range of payout parameterizations (see Appendix A for a complete report of experimental results).

Compared Algorithms. Each simulation compares the performance of four variants of Thompson Sampling: (1) TS is the TS bandit algorithm that attempts to maximize $E[y|do(x)]$, and does not condition on intent; (2) TS^I is a contextual TS player that uses intent as a context; (3) TS^{RDC} is the approach produced by [2], which conditions on intent and employs observational data, but does not incorporate experimental data nor exploit the relationship between data types detailed in the previous section; (4) TS^{RDC+} follows the algorithm described above.

Evaluation. Each algorithm’s performance is evaluated using two standard metrics: (1) the probability of optimal arm choice and (2) cumulative regret, both as a function of t averaged across all N Monte Carlo Simulations. However, unlike in the traditional MAB literature, we compare each algorithm’s choice to the optimal choice of an omniscient oracle that knows the value of any UCs in any given round of any MC repetition (indicated as $x_{n,t}^*$). Formally, for all $0 < t < T$ we evaluate (1) as $\frac{1}{N} \sum_n 1(x_{n,t}^* = x_{n,t})$ and (2) as $\frac{1}{N} \sum_n \sum_i^t y_{n,i} - E[Y_{x_{n,t}^*}|u_{n,i}]$.

Experiment 1: “Greedy Casino.” The Greedy Casino parameterization as, described in Table 1, exemplifies the scenario where all arms are both observationally equivalent ($E[Y|x] = E[Y|x'], \forall x, x'$) and experimentally equivalent ($E[Y|do(x)] = E[Y|do(x'), \forall x, x'$), but distinguishable within intent conditions ($E[Y_x|x']$). In this reward parameterization, TS^{RDC+} experienced significantly less regret ($M = 42.23$) than its chief competitor, TS^{RDC} , ($M = 65.04$), $t(1998) = 13.25, p < .001$.

⁷The parameters for these distributions are decided by the agent’s counterfactual data history (see Figure 4), including contributions from observational data for cells in which action and intent agree.

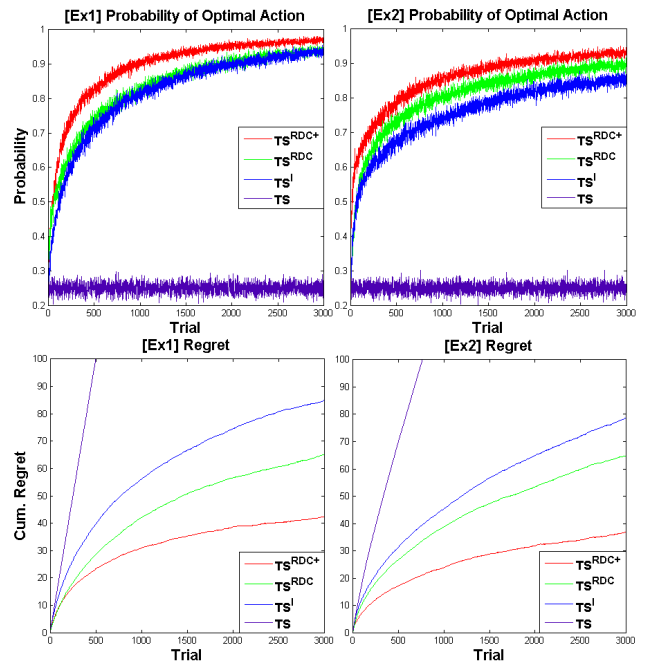


Figure 5: Plots of TS variant performances in the Greedy Casino [Ex1] and Paradoxical Switching [Ex2] scenarios.

Experiment 2: “Paradoxical Switching.” The Paradoxical Switching parameterization (see Table 2) exemplifies a curious scenario wherein $E[Y_{x_1}] = 0.5 > E[Y_{x'}], \forall x' \neq x_1$, but for which x_1 is the optimal arm choice in only one intent condition ($I = x_1$). Agents unempowered by RDC will face a paradox in that the arm with the highest experimental payout is not always optimal. Again, TS^{RDC+} experienced significantly less regret ($M = 36.91$) than its chief competitor, TS^{RDC} , ($M = 64.70$), $t(1998) = 22.43, p < .001$.

The accelerated learning enjoyed by $RDC+$ is not localized to these parameter choices alone. In Appendix A (Supplemental Material), we show that TS^{RDC+} consistently experiences significantly less regret than its competitors across a wide range of parametrizations that encode qualitatively different relationships between the observational, experimental, and counterfactual distributions.

6. DISCUSSION

This work fits into the broader literature of autonomous agent design that will require more generalized actors to distinguish between the sources, semantics, and personalized decision-making implications of the data that they collect. In particular, the Regret Decision Criterion (and counterfactual data, more generally) serves as a meta-decision-making tool (i.e., a tool used by an agent to reflect upon its own decision-making processes) for evaluating policy efficacy and how well an agent’s heuristics generalize across different environments.

Various meta-analytic approaches for tuning AI systems have begun to appear in the literature, not only corroborating merit for such approaches, but also suggesting future applications of RDC. For instance, upcoming studies

(a)	$D = 0$		$D = 1$	
	$B = 0$	$B = 1$	$B = 0$	$B = 1$
$E[y_1 X, B, D]$				
$X = 0$	*0.90	0.20	0.45	0.45
$X = 1$	0.30	*0.40	0.50	0.40
$X = 2$	0.10	0.35	*0.60	0.35
$X = 3$	0.10	0.10	0.30	*0.60

(b)	$E[y_1 X]$	$E[y_1 do(X)]$
$X = 0$	0.90	0.50
$X = 1$	0.40	0.40
$X = 2$	0.60	0.35
$X = 3$	0.60	0.20

Table 2: (a) Paradoxical switching scenario payout rates decided by reactive slot machines as a function of arm choice X , sobriety D , and machine conspicuosity B . Players’ natural arm choices under D, B are indicated by asterisks. (b) Payout rates according to the observational, $E[y_1|X]$, and experimental $E[y_1|do(X)]$, distributions, where $Y = y_1$ represents winning (shown in the table).

have sought to increase the longevity or generalization capacities of legacy information systems as aging applications must adapt to changes in their environments that may have been overlooked during development [8]. Others have begun exploring “deep-meta-reinforcement learning” for recurrent neural networks in reinforcement learning domains [17].

More akin to the present work, other research has begun to address the data-fusion problem from a causal perspective, though without a focus on the effects of UCs on personalized decision-making as presented herein [7]. We forecast that active reinforcement learners of the adjacent future will place greater emphases on concerting data-fusion with personalized decision-making in the face of UCs, and that the present work begins to bridge these previously separate topics.

7. CONCLUSION

The present work addresses the challenges faced by on-line learning agents that gain access to increasingly diverse, and qualitatively different, sources of information implicit in real-world environments, and how these sources can be meaningfully synthesized to accelerate learning. This data-fusion problem is complicated by the presence of unobserved confounders (UCs), whose identities and influences are unknown to the actor. In response, we present a novel, counterfactual method by which online agents may combine their observations, experiments, and counterfactual (i.e., personalized) experiences to more quickly learn about their environments, even in the presence of UCs. This method is facilitated by our proof that agents practicing intent-specific experimentation can measure counterfactual quantities that were previously thought to be unidentifiable in the absence of a fully-specified model. We illustrate the efficacy of this approach in the Multi-Armed Bandit problem with Unobserved Confounders (MABUC), and demonstrate how a traditional bandit algorithm (Thompson Sampling) may be improved by its application. Simulations demonstrate that our data-fusion approach generalizes across reward parameterizations and results in significantly less regret (in some cases, as much as half) than other competitive MABUC algorithms.

REFERENCES

- [1] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.
- [2] E. Bareinboim, A. Forney, and J. Pearl. Bandits with unobserved confounders: A causal approach. In *Advances in Neural Information Processing Systems*, pages 1342–1350, 2015.
- [3] E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 2016.
- [4] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5:1–122, 2012.
- [5] N. R. Council et al. *Frontiers in massive data analysis*. National Academies Press, 2013.
- [6] E. Even-Dar, S. Mannor, and Y. Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *J. Mach. Learn. Res.*, 7:1079–1105, Dec. 2006.
- [7] P. Geiger, K. Hofmann, and B. Schölkopf. Experimental and causal view on information integration in autonomous agents. *arXiv preprint arXiv:1606.04250*, 2016.
- [8] A. N. Gorban, I. Romanenko, R. Burton, and I. Y. Tyukin. One-Trial Correction of Legacy AI Systems and Stochastic Separation Theorems. *arXiv preprint arXiv:1610.00494v2*, 2016.
- [9] M. K. Ho, M. Littman, J. MacGlashan, F. Cushman, J. Austerweil, and J. L. Austerweil. Showing versus doing: Teaching by demonstration. In *Advances In Neural Information Processing Systems*, pages 3027–3035, 2016.
- [10] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4 – 22, 1985.
- [11] X.-L. Meng. A trio of inference problems that could win you a nobel prize in statistics (if you help fund it). *Past, Present, and Future of Statistical Science*, 2014.
- [12] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2000. Second ed., 2009.
- [13] J. Pearl, M. Glymour, and N. P. Jewell. *Causal Inference in Statistics: A Primer*. Wiley, 2016.
- [14] H. Robbins. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.*, 58(5):527–535, 09 1952.
- [15] S. L. Scott. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658, 2010.
- [16] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*, pages 141–162. Springer, 1975.
- [17] J. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick. Learning to Reinforcement Learn. *arXiv preprint arXiv:1611.05763v1*, 2016.

- [18] J. Zhang and E. Bareinboim. Markov decision processes with unobserved confounders: A causal approach. Technical Report R-23, Purdue AI Lab, 2016.