

Exploiting Action Impact Regularity and Exogenous State Variables for Offline Reinforcement Learning

Vincent Liu
James R. Wright
Martha White

VLIU1@UALBERTA.CA
JAMES.WRIGHT@UALBERTA.CA
WHITEM@UALBERTA.CA

*University of Alberta and Alberta Machine Intelligence Institute (Amii)
Edmonton, Alberta, Canada*

Abstract

Offline reinforcement learning—learning a policy from a batch of data—is known to be hard for general MDPs. These results motivate the need to look at specific classes of MDPs where offline reinforcement learning might be feasible. In this work, we explore a restricted class of MDPs to obtain guarantees for offline reinforcement learning. The key property, which we call Action Impact Regularity (AIR), is that actions primarily impact a part of the state (an endogenous component) and have limited impact on the remaining part of the state (an exogenous component). AIR is a strong assumption, but it nonetheless holds in a number of real-world domains including financial markets. We discuss algorithms that exploit the AIR property, and provide a theoretical analysis for an algorithm based on Fitted-Q Iteration. Finally, we demonstrate that the algorithm outperforms existing offline reinforcement learning algorithms across different data collection policies in simulated and real world environments where the regularity holds.

1. Introduction

Offline reinforcement learning (RL) involves using a previously collected static dataset, without any online interaction, to learn an output policy. This problem setting is important for a variety of real world problems where learning online can be dangerous, such as for self-driving cars, or when building a good simulator is difficult or costly, such as for healthcare. It is also a useful setting for applications where there is a large amount of data available, such as dynamic search advertising.

A challenge in offline RL is that the quality of the output policy can be highly dependent on the data. Most obviously, the data might not cover some parts of the environment, resulting in two issues. The first is that the learned policy, when executed in the environment, is likely to deviate from the behavior that generated its training data and reach a state-action pair that was unseen in the dataset. For these unseen state-action pairs, the algorithm has no information about how to choose a good action. The second issue is that if the dataset does not contain transitions in the high-reward regions of the state-action space, it may be impossible for any algorithm to return a good policy. One can easily construct a family of MDPs with missing data such that no algorithm can identify the MDP and suffer a large suboptimality gap (Chen & Jiang, 2019).

The works that provide guarantees on the suboptimality of the output policy usually rely on strong assumptions about good data coverage and mild distribution shift. The theoretical results are for methods based on approximate value iteration (AVI) and approximate

policy iteration (API), with results showing the output policy is close to an optimal policy (Farahmand, Szepesvári, & Munos, 2010; Munos, 2003, 2005, 2007). They assume a small concentration coefficient, which is the ratio between the state-action distribution induced by any policy and the data distribution (Munos, 2003). However, the concentration coefficient can be very large or even infinite in practice, so assuming a small concentration coefficient can be unrealistic for many real world problems. Different measures of distribution shift are also used, for example, Yin, Bai, and Wang (2021) assume the visitation distribution of the least occupied state-action pair is greater than zero.

To avoid making strong assumptions on the concentration coefficient, several works consider constraining divergence between the behavior policy and the output policy on the policy improvement step for API algorithms. The constraints can be enforced either as direct policy constraints or by a penalty added to the value function (Levine, Kumar, Tucker, & Fu, 2020; Wu, Tucker, & Nachum, 2019). Another approach is to constrain the policy set such that it only chooses actions or state-action pairs with sufficient data coverage when applying updates for AVI algorithms (Kumar, Fu, Soh, Tucker, & Levine, 2019; Liu, Swaminathan, Agarwal, & Brunskill, 2020). However, these methods only work when the data collection policy covers an optimal policy (see our discussion in Section 3), which can be difficult or impossible to guarantee.

Another direction has been to assume pessimistic values for unknown state-action pairs, to encourage the agent to learn an improved policy that stays within the parts of the space covered by the data. CQL (Kumar, Zhou, Tucker, & Levine, 2020) penalizes the values for out-of-distribution actions and learns a lower bound of the value estimates. A related idea is to constrain the bootstrap target to avoid out-of-distribution actions, introduced first in the BCQ algorithm (Fujimoto, Meger, & Precup, 2019) with practical improvements given by IQL (Kostrikov, Nair, & Levine, 2022). MOREL (Kidambi, Rajeswaran, Netrapalli, & Joachims, 2020) learns a model and an unknown state-action detector to partition states similar to the R-Max algorithm (Brafman & Tenenbholz, 2002), but then uses the principle of pessimism for these unknown states rather than optimism. Safe policy improvement methods (Thomas, Theodorou, & Ghavamzadeh, 2015; Larocche, Trichelair, & Des Combes, 2019) rely on a high-confidence lower bound on the output policy performance, performing policy improvement only when the performance is higher than a threshold.

In practice, the results of pessimistic approaches are mixed. Some have been shown to be effective on the D4RL dataset (Fu, Kumar, Nachum, Tucker, & Levine, 2020). Other results, however, show methods can be too conservative and fail drastically when the behavior policy is not near-optimal (Kumar et al., 2020; Liu et al., 2020), as we reaffirm in our experiments. Further, actually using these pessimistic methods can be difficult, as their hyper-parameters are not easy to tune in the offline setting (Wu et al., 2019) and some methods require an estimate of the behavior policy or data distribution (Kumar et al., 2019; Liu et al., 2020).

Intuitively, however, there are settings where offline RL should be effective. Consider a trading agent in a stock market. A policy that merely observes stock prices and volumes without buying or selling any shares provides useful information about the environment. For this collected dataset, an offline agent can counter-factually reason about the utility of many different actions as demonstrated in Figure 1, because its actions have limited impact on the prices and volumes. Such MDPs, which are called Exogenous MDPs, have states that separate into exogenous states (stock price), not impacted by actions, and endogenous

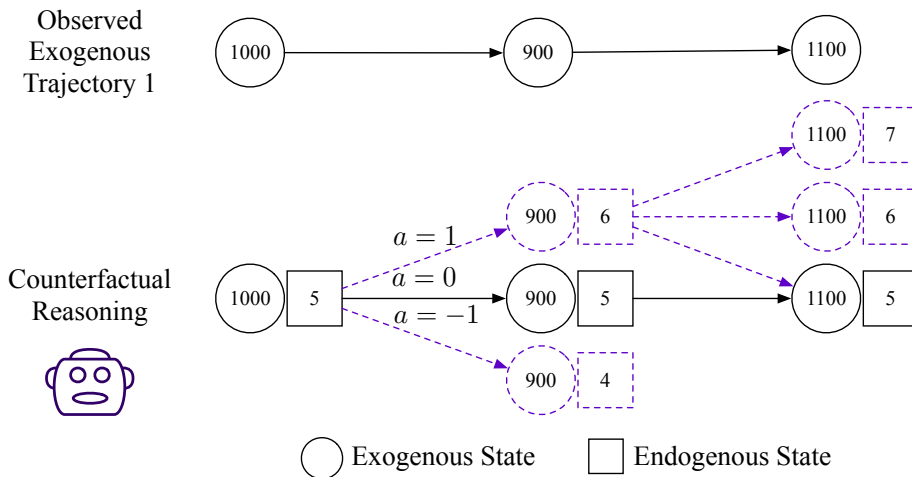


Figure 1: In the stock market example, the exogenous state corresponds to the stock price, the endogenous state corresponds to number of shares the agent has, and the action corresponds to the number of share to buy or sell at each time step. Given an observed exogenous trajectory, the agent can counter-factually reason about the outcomes of different actions and endogenous state.

states (number of shares owned by the agent). The structure in Exogenous MDPs has been used in online RL to learn more efficiently (Dietterich, Trimponias, & Chen, 2018).

This exogenous structure, however, has yet to be formally investigated for offline RL, though it is likely already being exploited in industry. Exploiting this structure is natural in applied financial applications, because datasets allow for alternative trajectories to be simulated, as described in the above example. One (unpublished) system uses RL and trajectory simulation for the optimal order execution problem (Burhani, Ding, Hernandez-Leal, Prince, Shi, & Szeto, 2020); it seems likely that there are other such systems in use. What has yet to be done, however, is to understand the theoretical properties of such algorithms, as well as potential algorithmic improvements.

In this paper, we first generalize the definition of exogenous MDPs, and formalize the action impact regularity (AIR) property. We say an MDP has the AIR property—or is ε -AIR—if the actions have a limited impact on exogenous dynamics, with the level of impact determined by $\varepsilon \geq 0$. This generalizes the previous definition, which required strict separation, corresponding to $\varepsilon = 0$. We develop both theory and algorithms for this more general setting, assuming access only to an offline dataset, an approximate (learned) endogenous model and the reward function.¹

We design an efficient algorithm, called FQI-AIR, to exploit the AIR property, that (1) does not require an estimate of the behavior policy or the data distribution, (2) has a straightforward approach to select hyperparameters using just the given data and (3) is much less sensitive to the quality of the data collection policy, if our assumptions hold. This algorithm is a simple extension of FQI, but is significantly more computationally efficient

1. We could assume an approximate instead of exact reward model. However, in most RL analysis, the error on the transition model is of a higher order than the error for reward models (for example, see Agarwal, Kakade, and Yang (2020)). For simplicity, it is often assumed the reward model is known.

than the trajectory simulation approach mentioned above and allows us to leverage and extend the existing theory for FQI. We bound the suboptimality of the output policy from FQI-AIR, in terms of ε and other standard terms (model errors and the inherent Bellman error, see Section 5.2). Importantly, in place of the concentration coefficient, we have a term that depends on the size of the endogenous state and number of actions; when the concentration coefficient is bounded and small, it is on the same order as this term.

We then conduct a comprehensive empirical study of FQI-AIR. We compare several algorithms in two simulated environments, across three different data collection policies, with varying offline dataset sizes, for $\varepsilon = 0$ (assumption perfectly satisfied) and a larger ε (assumption somewhat violated). FQI-AIR significantly outperforms the offline RL algorithms that do not leverage the AIR property—including FQI, MBS-QI, CQL and IQL; this outcome is expected, but nonetheless verifies that exploiting the AIR property, when appropriate, can have a big benefit. We show that these conclusions extend to two environments based on real-world datasets (for bitcoin trading and for controlling battery usage in a hybrid car). An important detail here is how hyperparameters are chosen. FQI-AIR can exploit AIR for policy evaluation, to automatically select hyperparameters. For the other algorithms, we do not have such an approach, and instead report idealized performance by picking hyperparameters based on performance in the environment.

Finally, these results all used the true endogenous model for FQI-AIR. We chose to do so partly because the endogenous model is known for certain AIR-MDPs (e.g., trading, inventory management) and partly to focus the investigation on the role of ε rather than model error. However, for certain AIR-MDPs, we will not have access to the true endogenous model. In our final experiment, we investigated the impact of using a learned endogenous model in the hybrid car environment, and show that FQI-AIR remains effective.

All of this is only possible because we make a strong assumption about the environment. However, given the hardness results in offline RL, we should acknowledge that we likely need to restrict the class of MDPs. This work is a step towards understanding for what classes of MDPs offline RL is feasible. At the same time, though we consider a restricted setting, it is by no means a trivial setting. There are many real-world examples where this regularity holds (as we discuss later in this work). This is doubly true given that our generalization provides some flexibility in violating the assumption: the regularity only needs to hold approximately rather than exactly. The algorithms and theory developed here can benefit these real-world applications now, by providing an approach that is well-designed and well-behaved with strong theoretical guarantees for their specific problem setting.

2. Problem Formulation

The agent-environment interaction is formalized as a finite horizon Markov decision process (MDP) $M = (\mathcal{S}, \mathcal{A}, P, r, H, \nu)$. \mathcal{S} is a set of states, and \mathcal{A} is an set of actions; for simplicity, we assume that both sets are finite. $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition probability where $\Delta(\mathcal{S})$ is the set of probability distributions on \mathcal{S} , and by slightly abusing notation, we will write $P(s, a, s')$ as the probability that the process will transition into state s' when in state s it takes action a . The function $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, r_{max}]$ gives the reward when taking action a in state s , where $r_{max} \in \mathbb{R}^+$. $H \in \mathbb{Z}^+$ is the planning horizon, and $\nu \in \Delta(\mathcal{S})$ the initial state distribution.

In the finite horizon setting, the policies are non-stationary. A non-stationary policy is a sequence of memoryless policies $(\pi_0, \dots, \pi_{H-1})$ where $\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. We assume that the set of states reachable at time step h , $\mathcal{S}_h \subset \mathcal{S}$, are disjoint, without loss of generality, because we can always define a new state space $\mathcal{S}' = \mathcal{S} \times [H-1]$ where $[n] := \{0, 1, 2, \dots, n\}$. Then, it is sufficient to consider stationary policies $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$.

Given a policy π , $h \in [H-1]$, and $(s, a) \in \mathcal{S} \times \mathcal{A}$, we define the value function and the action-value function as $v_h^\pi(s) := \mathbb{E}^\pi[\sum_{t=h}^{H-1} r(S_t, A_t) | S_h = s]$ and $q_h^\pi(s, a) := \mathbb{E}^\pi[\sum_{t=h}^{H-1} r(S_t, A_t) | S_h = s, A_h = a]$ where the expectation is with respect to \mathbb{P}_M^π (we may drop the subscript M when it is clear from the context). \mathbb{P}^π is the probability measure on the random element in $(\mathcal{S} \times \mathcal{A})^H$ induced by the policy π and the MDP such that, for the trajectory of state-action pairs $(S_0, A_0, \dots, S_{H-1}, A_{H-1})$, we have $\mathbb{P}^\pi(S_0 = s) = \nu(s)$, $\mathbb{P}^\pi(A_t = a | S_0, A_0, \dots, S_t) = \pi(a | S_t)$, and $\mathbb{P}^\pi(S_{t+1} = s' | S_0, A_0, \dots, S_t, A_t) = P(S_t, A_t, s')$ for $t \geq 0$ (Lattimore & Szepesvári, 2020). The optimal value function is defined by $v_h^*(s) := \sup_\pi v_h^\pi(s)$, and the Bellman operator is defined by

$$(\mathcal{T}q_h)(s, a) = r(s, a) + \sum_{s' \in \mathcal{S}} P(s, a, s') \max_{a' \in \mathcal{A}} q_h(s', a').$$

In the batch setting, we are given a fixed set of transitions D with samples drawn from a data distribution μ over $(\mathcal{S} \times \mathcal{A})$. In this paper, we consider the setting where the data is collected by a data collection policy π_b . That is, D consists of N trajectories $(S_0^{(i)}, A_0^{(i)}, \dots, S_{H-1}^{(i)}, A_{H-1}^{(i)})$ induced by the the interaction of the policy π_b and MDP M .

A representative algorithm for the batch setting is *Fitted Q Iteration* (FQI) (Ernst, Geurts, & Wehenkel, 2005). In the finite horizon setting, FQI learns an action-value function for each time step, $q_0, \dots, q_{H-1} \in \mathcal{F}$ where $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ is the value function class. The algorithm is defined recursively from the end of the episode: for each time step h from $H-1$ to 0, $q_h = \arg \min_{q \in \mathcal{F}} \|q - \hat{\mathcal{T}}q_{h+1}\|_2^2$ where $\hat{\mathcal{T}}$ is defined by replacing expectations with sample averages for the Bellman operator \mathcal{T} and $q_H = 0$. The output policy is obtained by greedifying according to these action-value functions.

3. Common Assumptions for Offline RL

In this section we discuss common assumptions used in offline RL. The most common assumptions typically concern properties of the data distribution and MDPs together: either that it sufficiently covers the set of possible transitions, or that it sufficiently covers a near-optimal policy. However, these assumptions are often impractical for many real world applications. Furthermore, recent results show that restriction on the data distribution alone are insufficient to obtain guarantees. This motivates considering realistic assumptions on the MDP alone, as we do in this work.

Sufficient coverage of the data distribution has primarily been quantified by *concentration coefficients* (Munos, 2003). Given a data distribution μ , the concentration coefficient C is defined to be the smallest value such that, for any policy π ,

$$\max_{h \in [H-1]} \max_{s \in \mathcal{S}_h, a \in \mathcal{A}_h} \frac{\mathbb{P}^\pi(S_h = s, A_h = a)}{\mu(s, a)} \leq C.$$

If $\mu(s, a) = 0$ for some (s, a) , then we define $C = \infty$ by convention. Several results bound the suboptimality of batch API and AVI algorithms in terms of the concentration coefficient (Chen & Jiang, 2019; Farahmand et al., 2010; Munos, 2003, 2007). For example, it has been shown that FQI outputs a near-optimal policy when C is small and the value function class is rich enough, where the upper bound on the suboptimality of the output policy scales linearly with \sqrt{C} .

However, in practice, the concentration coefficient can be very large or even infinite. For example, if the data collection policy is not well-randomized or exploratory—often the case in practice—then the concentration coefficient is infinite due to missing some state-action pairs. Munos (2007) provides some intuition about the size of the concentration coefficient. Suppose that the data distribution is uniform (e.g., $\mu(s, a) = 1/|\mathcal{S}||\mathcal{A}|$) and the environment transition probability is less uniform, that is, there exists some policies such that the visitation distribution concentrates on a single state-action pair (e.g., $\mathbb{P}^\pi(S_h = s, A_h = a) = 1$ for some s, a and h), then the concentration coefficient can be as large as the number of state-action pairs

Another direction has been to consider approaches where the data covers a near-optimal policy. The key idea behind these methods is to restrict the policy to choose actions that have sufficient data coverage, which is effective if the given data has near-optimal action selection. For example, BCQ (Fujimoto et al., 2019) and BEAR (Kumar et al., 2019) only bootstrap values from actions a if the probability $\pi_b(a|s)$ is above a threshold b . MBS-QI (Liu et al., 2020) extends this to consider state-action probabilities, only bootstrapping from state-action pairs (s, a) when $\mu(s, a)$ is above a threshold. The algorithm is modified from FQI by replacing the bootstrap value $q_h(s, a)$ by $\tilde{q}_h(s, a) := \mathbb{I}\{\mu(s, a) \geq b\}q_h(s, a)$ and the policy is greedy with respect to $\tilde{q}_h(s, a)$. If a state-action pair does not have sufficient data coverage, its value is zero. They show that MBS-QI outputs a near-optimal policy if $\mathbb{P}^{\pi^*}(\mu(S_h, A_h) < b)$ is small for all $h \in [H-1]$. That is, the data provides sufficient coverage for state-action pairs visited under an optimal policy π^* .

Though potentially less stringent than having a small concentration coefficient, this assumption can be impractical. We may be able to satisfy this assumption in simulated environments, such as those in our experiments; in the real world, though, if we have a simulator we are unlikely to use offline RL. For many real world domains, optimal policies are unknown. In fact, one of the primary purposes of using offline RL is to get (significantly) improved policies. It is also hard to carefully design a data collection policy to cover an unknown optimal policy, making it difficult to even check whether this assumption holds.

Finally, some recent negative results suggest it is not sufficient to have a good data distribution alone, and that it will be necessary to also make assumptions on the MDP. In particular, Chen and Jiang (2019) showed that if we do not make assumptions on the MDP dynamics, no algorithm can achieve a polynomial sample complexity to return a near-optimal policy, even when the algorithm can choose any data distribution μ . Wang, Foster, and Kakade (2021) provide an exponential lower bound for the sample complexity of off-policy policy evaluation and optimization algorithms with q^π -realizable linear function class, even when the data distribution induces a well-conditioned covariance matrix. Zanette (2021) provide an example where offline RL is exponentially harder than online RL, even with the best data distribution, q^π -realizable function class and assuming the exact feedback is observed for each sample in the dataset. Xiao, Lee, Dai, Schuurmans,

and Szepesvari (2022) provide an exponential lower bound for the sample complexity of obtaining nearly-optimal policies when the data is obtained by following a data collection policy. These results are consistent with the above, since achieving a small concentration coefficient implicitly places assumptions on the MDP.

The main message from these negative results is that a good data distribution alone is not sufficient. We need to investigate realistic problem-dependent assumptions for MDPs. In the remainder of this work, we explore a restricted class of MDPs, for which we can obtain much stronger guarantees when learning offline, without stringent requirements on data collection.

4. Action Impact Regularity

Actions play an important role for the exponential lower bound constructions cited in the last section. They use tree structures where different actions lead to different subtrees and hence different sequence of futures states and rewards. A class of MDPs that do not suffer from these lower bounds are those where actions do not have such strong impact on the future states and rewards. In this section, we introduce the Action Impact Regularity (AIR) property, a property of the MDP which allows for more effective offline RL. The state is partitioned into an exogenous and endogenous component, and the property reflects that the agent’s actions primarily impact the endogenous state with limited influence on the exogenous state. We first provide the formal definition and assumptions we leverage to design a practical offline RL algorithm and then discuss when these assumptions are likely to be satisfied.

4.1 Formal Definition and Assumptions

We use the standard state decomposition from Exogenous MDPs (McGregor, Houtman, Montgomery, Metoyer, & Dietterich, 2017; Dietterich et al., 2018). We assume the state space is $\mathcal{S} = \mathcal{S}^{\text{exo}} \times \mathcal{S}^{\text{end}}$ where \mathcal{S}^{exo} is the exogenous variable and \mathcal{S}^{end} is the endogenous variable. The transition dynamics are $P^{\text{exo}} : \mathcal{S}^{\text{exo}} \times \mathcal{A} \rightarrow \Delta(\mathcal{S}^{\text{exo}})$ and $P^{\text{end}} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S}^{\text{end}})$ for exogenous and endogenous variable respectively. The transition probability from a state $s_1 = (s_1^{\text{exo}}, s_1^{\text{end}})$ to another state $s_2 = (s_2^{\text{exo}}, s_2^{\text{end}})$ is $P(s_1, a, s_2) = P^{\text{exo}}(s_1^{\text{exo}}, a, s_2^{\text{exo}})P^{\text{end}}(s_1, a, s_2^{\text{end}})$.

Definition 1 (The AIR Property). An MDP is ε -AIR if $\mathcal{S} = \mathcal{S}^{\text{exo}} \times \mathcal{S}^{\text{end}}$, and for any actions $a, a' \in \mathcal{A}$, the next exogenous variable distribution is similar if either action a or a' is taken. That is, for each state $s \in \mathcal{S}$,

$$D_{TV}(P^{\text{exo}}(s^{\text{exo}}, a), P^{\text{exo}}(s^{\text{exo}}, a')) \leq \varepsilon$$

where D_{TV} is the total variation distance between two probability distributions on \mathcal{S}^{exo} . For discrete spaces, the total variation distance is $D_{TV}(P, P') = \frac{1}{2}\|P - P'\|_1$ (ℓ_1 norm).

We define the AIR-MDP such that the property holds for all exogenous state-action pairs. If the property does not hold for one of the exogenous state-action pairs, then one can design an adversarial MDP that hides all difficulties in this single exogenous state-action pair and assuming the properties hold for all but one pair would be useless (Jiang, 2018).

Access to an (approximate) endogenous model is critical to exploit the AIR property, and is a fundamental component of our algorithm. To be precise, we make the following assumption in this paper.

Assumption 1 (AIR with an Approximate Endogenous Model). We assume that the MDP is ε_{air} -AIR and that we have the reward model $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, r_{max}]$ and an approximate endogenous model $\hat{P}^{\text{end}} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S}^{\text{end}})$ such that $D_{TV}(P^{\text{end}}(s, a), \hat{P}^{\text{end}}(s, a)) \leq \varepsilon_p$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.

As mentioned in the introduction, it is common to assume that only the transition dynamics are approximated. Moreover, similar to Definition 1, we need the error on the approximate model to hold uniformly. Finally, the above assumption implicitly assumes that the separation between exogenous and endogenous state is given to us. More generally, the separation could be identified or learned by the agent, as has been done for contingency-aware RL agents (Bellemare, Veness, & Bowling, 2012) and wireless networks (Dietterich et al., 2018). Because there are many settings where the separation is clear, we focus on this more clear case first where the separation is known.

4.2 When Are These Assumptions Satisfied?

Many real-world problems can be formulated as ε -AIR MDPs. Further, for many of these environments, the separation between exogenous and endogenous state is clear, and we either know or can reasonably approximate the endogenous model. In this section, we go through several concrete examples.

We can first return to our stock trading example, from the introduction. The exogenous component is the market information (stock prices and volumes) and the endogenous component is the number of stock shares owned by the agent. The agent’s actions influence their own number of shares, but as an individual trader, have limited impact on stock prices. Using a dataset of stock prices over time allows the agent to reason counterfactually about the impact of many possible trajectories of actions (buying/selling) on its shares (endogenous state) and profits (reward).

There are many settings where the agent has a limited impact on a part of the state. The optimal order execution problem is a task to sell M shares of a stock within H steps; the goal is to maximize the profit. The problem can be formulated as an MDP where the exogenous variable is the stock price and endogenous variable is the number of shares left to sell. It is common to assume infinite market liquidity (Nevmyvaka, Feng, & Kearns, 2006) or that actions have a small impact on the stock price (Abernethy & Kale, 2013; Bertsimas & Lo, 1998); this corresponds to assuming the AIR property.

Another example is the secretary problem (Freeman, 1983), which a family of problems that can often be used to model real-world application (Babaioff, Immorlica, Kempe, & Kleinberg, 2007; Goldstein, McAfee, Suri, & Wright, 2020). The goal for the agent is to hire the best secretary out of H , interviewed in random order. After each interview, they have to decide if they will hire that applicant, or wait to see a potentially better applicant in the future. The problem can be formulated as a 0-AIR MDP where the endogenous variable is a binary variable indicating whether we have chosen to stop or not.

Other examples include those where the agent only influences energy efficiency, such as in the hybrid vehicle problem (Shahamiri, 2008; Lian, Peng, Wu, Tan, & Zhang, 2020) and

electric vehicle charging problem (Abdullah, Gastli, & Ben-Brahim, 2021). In the former problem, the agent controls the vehicle to use either the gas engine or the electrical motor at each time step, with the goal to minimize gas consumption; its actions do not impact the driver’s behavior. In the latter problem, the agent controls the charging schedule of an electric vehicle to minimize costs; its actions do not impact electricity cost.

In some settings we can even restrict the action set or policy set to make the MDP ϵ -AIR. For example, if we know that selling M shares hardly impacts the markets, we can restrict the action space to selling less than or equal to M shares. In the hybrid vehicle example, if the driver can see which mode is used, we can restrict the policy set to only switch actions periodically to minimize distractions for the driver.

In these problems with AIR, we often know the reward and transitions for the endogenous variables, or have a good approximation. For the optimal order execution problem, the reward is simply the selling price times the number of shares sold minus transaction costs, and the transition probability for P^{end} is the inventory level minus the number of shares sold. In other applications, we may be able to use domain knowledge to build an accurate model for the endogenous dynamics. For the hybrid vehicle, we can use domain knowledge to calculate how much gas would be used for a given acceleration. Such information about the dynamics of the system can be simpler for engineers to specify, than (unknown) behavior of different drivers and environment conditions. Our theoretical results will include a term for the error in the endogenous model, but it is reasonable to assume that for many settings we can get that error to be relatively low, particularly in comparison to the error we might get if trying to model the exogenous state.

4.3 Connections to the Literature on Exogenous MDPs

AIR MDPs can be viewed as an extension of Exogenous MDPs. (1) We allow the action to have small impact on the environmental state, while the action has no impact on the exogenous state in Exogenous MDPs. (2) We do not assume the reward can be decomposed additively to an exogenous reward and an endogenous reward (Dietterich et al., 2018) nor factor into a sum over each exogenous state variable (Chitnis & Lozano-Pérez, 2020). For this previous definition of Exogenous MDPs, the focus was on identifying and removing the exogenous state/noise so that the learning problem could be solved more efficiently (Dietterich et al., 2018; Efroni, Misra, Krishnamurthy, Agarwal, & Langford, 2022), thus the focus on reward decomposition. Our focus is offline learning where we want to exploit the known structure to enable counterfactual reasoning and avoid data coverage issues.

5. Offline Policy Optimization for AIR MDPs

In this section, we discuss several offline algorithms that exploit the AIR property for policy optimization. We then theoretically analyze an FQI-based algorithm, characterizing the performance of its outputted policy.

5.1 Algorithms for AIR MDPs

Two standard classes of algorithms in offline RL are model-based algorithms—that learn a model from the offline dataset and then use dynamic programming—and model-free algo-

rithms like fitted Q-iteration (FQI). These two approaches can be tailored to our setting with AIR MDPs, as we described below. There is, however, an even more basic approach in our offline RL setting using trajectory simulation, that has previously been used (Burhani et al., 2020). We start by describing this simpler approach, and then the modified model-based and FQI approaches.

A natural approach is to reuse trajectories in the dataset to simulate alternative trajectories for an online RL algorithm. For each episode, a random trajectory is selected from the dataset. The online RL algorithm—such as an actor-critic method or a Q-learning agent—takes actions and deterministically transitions to the next exogenous state in the trajectory. The approximate endogenous and reward model are used to sample the next endogenous variable and reward. With such a trajectory simulator, we can run any online reinforcement learning algorithm to find a good policy for the simulator.

This approach, however, does not exploit the fact that the agent is actually learning offline. The online RL algorithm cannot simply query the model for any state and action, and needs a good exploration strategy to find a reasonable policy. There are fewer theoretical guarantees for such online RL algorithms, and arguably more open questions about their properties than DP-based algorithms and fitted value iteration algorithms.

A more explicit model-based approach is to learn the exogenous model from data, to obtain a complete transition and reward model, and use any planning approach. The transition model for exogenous states can be constructed as if the action has no impact. With the model, we can use any query-efficient planning algorithm to find a good policy for the model. Because actions have only small impact in the true MDP, we can learn an accurate exogenous model even if we do not have full data coverage.

More precisely, recall the offline data is randomly generated by running π_b on M , that is, $D = \{(S_0^{(i)}, A_0^{(i)}, \dots, S_{H-1}^{(i)}, A_{H-1}^{(i)})\}_{i=1}^N$ sampled according to the probability measure $\mathbb{P}_M^{\pi_b}$. The pertinent part is the transitions between exogenous variables, so we define $D_{\text{exo}} = \{(S_0^{(i)}, S_1^{(i),\text{exo}}, \dots, S_{H-1}^{(i),\text{exo}})\}_{i=1}^N$. The model-based approach constructs an empirical MDP $M_D = (\mathcal{S}, \mathcal{A}, \hat{P}, r, H, \hat{v})$. For the tabular setting we have $\hat{v}(s) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(S_0^{(i)} = s)$, and

$$\hat{P}^{\text{exo}}(s_h^{\text{exo}}, a, s_{h+1}^{\text{exo}}) = \frac{\sum_{i=1}^N \mathbb{I}(S_h^{(i),\text{exo}} = s_h^{\text{exo}}, S_{h+1}^{(i),\text{exo}} = s_{h+1}^{\text{exo}})}{\sum_{i=1}^N \mathbb{I}(S_h^{(i),\text{exo}} = s_h^{\text{exo}})}$$

for all $a \in \mathcal{A}$. Exogenous variables not seen in the data are not reachable, and so can either be omitted from \hat{P}^{exo} or set to self-loop. For large or continuous state spaces, we can learn $p(s_{h+1}^{\text{exo}} | s_h^{\text{exo}})$ using any conditional distribution learning algorithm, and set $\hat{P}^{\text{exo}}(s_h^{\text{exo}}, a, s_{h+1}^{\text{exo}}) = p(s_{h+1}^{\text{exo}} | s_h^{\text{exo}})$ for all $a \in \mathcal{A}$.

For large or continuous states spaces, however, learning such a model and planning can be impractical. Learning an accurate exogenous model might be difficult if the exogenous transition is complex or the exogenous state is high-dimensional. Further, it is not possible to sweep through all states during planning. Smarter approximate dynamic programming algorithms need to be used, but even these can be quite computationally costly.

A reasonable alternative is FQI, which approximates value iteration without the need to learn a model. Our FQI algorithm that exploits the AIR property is described in Algorithm 1, which we call FQI-AIR. The algorithm simulates all actions from a state, and assumes it transitions to the exogenous state observed in the dataset. The reward and

Algorithm 1 FQI-AIR

 Input: dataset D , value function class \mathcal{F} , \hat{P}^{end} , r

 Let $q_H = 0$, $D_{H-1} = \emptyset$, ..., $D_1 = \emptyset$
for $h = H - 1, \dots, 0$ **do**

 For all $i \in \{1, \dots, N\}$, all $s_h^{\text{end}} \in \mathcal{S}^{\text{end}}$, all $a \in \mathcal{A}$

 Sample $s'^{\text{end}} \sim \hat{P}^{\text{end}}(s_h^{(i),\text{exo}}, s_h^{\text{end}}, a)$, compute target

$$t = r(s_h^{(i),\text{exo}}, s_h^{\text{end}}, a) + \max_{a' \in \mathcal{A}} q_{h+1}(s_{h+1}^{(i),\text{exo}}, s'^{\text{end}}, a')$$

 Add (synthetic) pair $((s_h^{(i),\text{exo}}, s_h^{\text{end}}, a), t)$ to D_h

 After generating D_h

$$q_h = \arg \min_{f \in \mathcal{F}} \sum_{(x,y) \in D_h} (f(x) - y)^2$$

$$\pi_h(s) = \arg \max_a q_h(s, a) \text{ for all } s \in \mathcal{S}_h$$

 Output: $\pi = (\pi_0, \dots, \pi_{H-1})$

endogenous state for each simulated action can be obtained using the reward model and approximate endogenous model. Even though the true MDP is not necessarily 0-AIR MDP, we will show in the analysis that as long as ε_{air} is small, the algorithm can return a nearly optimal policy in the true MDP. This algorithm, although simple, enjoys theoretical guarantees without making assumptions on the concentration coefficient, and can be much more computationally efficient than trajectory simulation methods.

Note that the computational cost scales with the size of \mathcal{S}^{end} and \mathcal{A} . When $|\mathcal{S}^{\text{end}}|$ or $|\mathcal{A}|$ is large, we can modify FQI-AIR to no longer use full sweeps. Instead, we can randomly sample from the endogenous state space and action. We include a practical implementation of FQI-AIR in Algorithm 2. For each exogenous state in the dataset, we sample an endogenous state and an action, and query the approximate model to obtain a target for FQI update. As a result, the computation can be independent of the size of \mathcal{S}^{end} and \mathcal{A} . However, for sample complexity, the performance loss of the algorithm would depend on the squared root of the size $|\mathcal{S}^{\text{end}}||\mathcal{A}|$, as shown in the next section.

5.2 Theoretical Analysis of FQI-AIR

First we need the following definitions. For a given MDP M , we define $J(\pi, M) := \mathbb{E}_M^\pi[R(\tau)]$ where $\tau = (S_0, A_0, \dots, S_{H-1}, A_{H-1})$ is a random element in $(\mathcal{S} \times \mathcal{A})^H$, the expectation \mathbb{E}_M^π is with respect to \mathbb{P}_M^π , and $R(\tau) = \sum_{h=0}^{H-1} r(S_h, A_h)$.

We also need the following assumption on the function approximation error. This is a common assumption to analyze approximate value iteration algorithms (Antos, Szepesvári, & Munos, 2006; Munos, 2007). Let $\tilde{\nu}_h(s_h^{\text{exo}}) = \mathbb{P}_M^{\pi_b}(S_h^{\text{exo}} = s_h^{\text{exo}})$ be the data distribution on

Algorithm 2 FQI-AIR for large state spaces

Input: dataset D , an approximate model \hat{P}^{end} , r , mini-batch size B , number of training iteration K , number of updates per iteration M

Initialize a Q function $q_\theta : \mathcal{S}^{\text{exo}} \times \mathcal{S}^{\text{end}} \times \mathcal{A} \times [H] \rightarrow \mathbb{R}$, parameterized by θ

$\bar{\theta} \leftarrow \theta$

for $k = 1, \dots, K$ **do**

for $m = 1, \dots, M$ **do**

 Sample a mini-batch of transitions $\{(s_j^{\text{exo}}, s_j^{\text{end}}, a_j, h_j, s_j^{\prime \text{exo}}, s_j^{\prime \text{end}})\}_{j=1}^B$ from D

 For all j , sample an endogenous state $\tilde{s}_j^{\text{end}} \in \mathcal{S}^{\text{end}}$ and an action $\tilde{a}_j \in \mathcal{A}$ randomly,

 sample $\tilde{s}_j^{\prime \text{end}} \sim \hat{P}^{\text{end}}(s_j^{\text{exo}}, \tilde{s}_j^{\text{end}}, \tilde{a}_j)$, and compute target

$$t_j = r(s_j^{\text{exo}}, \tilde{s}_j^{\text{end}}, \tilde{a}_j) + \max_{a' \in \mathcal{A}} q_{\bar{\theta}}(s_j^{\prime \text{exo}}, \tilde{s}_j^{\prime \text{end}}, a', h_j + 1)$$

 Compute the mini-batch loss $L(\theta) = \sum_{j=1}^B (q_\theta(s_j^{\text{exo}}, \tilde{s}_j^{\text{end}}, \tilde{a}_j, h_j) - t_j)^2$

 Update θ to reduce $L(\theta)$

$\bar{\theta} \leftarrow \theta$

Output: the greedy policy with respect to q_θ

\mathcal{S}^{exo} at horizon h . Given a probability measure ν_h on \mathcal{S}^{exo} and $p \in [1, \infty)$, define the norm

$$\|q\|_{p, \nu_h}^p = \sum_{s^{\text{exo}} \in \mathcal{S}^{\text{exo}}} \sum_{s^{\text{end}} \in \mathcal{S}^{\text{end}}} \sum_{a \in \mathcal{A}} \frac{\nu_h(s^{\text{exo}})}{|\mathcal{S}^{\text{end}}||\mathcal{A}|} |q(s^{\text{exo}}, s^{\text{end}}, a)|^p.$$

Assumption 2. Assume the function class \mathcal{F} is finite and the inherent Bellman error is bounded by ε_{apx} , that is,

$$\varepsilon_{\text{apx}} = \max_{h \in [H]} \max_{f' \in \mathcal{F}} \min_{f \in \mathcal{F}} \|\mathcal{T}f' - f\|_{2, \nu_h}^2.$$

We assume the function class is finite for simplicity, which is common in many offline RL papers (Chen & Jiang, 2019). If the function class is not finite but has a bounded complexity measure, we can derive similar results by replacing the size of the function class with the complexity measure. For example, Duan, Jin, and Li (2021) analyze FQI with the Rademacher complexity. Since studying the complexity measure is not a critical point for our paper, we decide to make the finite function class assumption.

Theorem 1. Under Assumption 1 and 2, let π_M^* be an optimal policy in M and π the output policy of FQI-AIR, then with probability at least $1 - \zeta$

$$J(\pi, M) \geq J(\pi_M^*, M) - 2v_{\text{max}}H(\varepsilon_{\text{air}} + \varepsilon_p) - (H + 1)H\sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{72v_{\text{max}}^2 \ln(H|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{\text{apx}} \right).$$

The bound on performance loss has three components: (1) a sampling error term which decreases with more trajectories; (2) the AIR parameter; and (3) an approximation error

term which depends on the function approximation used. The result implies that as long as we have a sufficient number of episodes, a good function approximation, and small ε_{air} , then the algorithm can find a nearly-optimal policy with high probability. For example, if $\varepsilon_{air}, \varepsilon_p$ and ε_{apx} are small enough, we only need $N = \tilde{O}(H^4 v_{max}^2 |\mathcal{S}^{\text{end}}| |\mathcal{A}| / \delta)$ trajectories, which is polynomial in H , to obtain a δ -optimal policy.

The proof can be found in the appendix. The key idea is to introduce a baseline MDP M_b that is 0-AIR, that approximates M which is actually ε_{air} -AIR. The baseline MDP $M_b = (\mathcal{S}, \mathcal{A}, \tilde{P}, r, H, \nu)$ has $\tilde{P}^{\text{exo}}(s_h^{\text{exo}}, a, s_{h+1}^{\text{exo}}) = \mathbb{P}_M^{\pi_b}(S_h^{\text{exo}} = s_h^{\text{exo}}, S_{h+1}^{\text{exo}} = s_{h+1}^{\text{exo}}) / \mathbb{P}_M^{\pi_b}(S_h^{\text{exo}} = s_h^{\text{exo}})$ and $\tilde{P}^{\text{end}}(s_h, a, s_h^{\text{end}}) = \hat{P}^{\text{end}}(s_h, a, s_h^{\text{end}})$. The transition probability for exogenous state does not depend on the action a taken, so M_b is 0-AIR. We show that FQI returns a good policy in M_b , and that good policies in M_b are also good in the true MDP M .

We can contrast this bound to others in offline RL. For FQI results that assume the concentration coefficient is bounded and small, the error bound has a term that scales with \sqrt{C} , which is on the same order as the term $\sqrt{|\mathcal{S}^{\text{end}}| |\mathcal{A}|}$ in our bound. We can get a similar bound by considering this restricted class of MDPs that are ε -AIR, without having to make any assumptions on the concentration coefficient. For settings where this assumption is appropriate—namely when the MDP is ε -AIR—this is a significant success, as we need not make these stringent conditions on data distributions.

6. Policy Evaluation for AIR MDPs

We can also exploit the AIR property, and access to the approximate endogenous model and reward model, to evaluate the value of a given policy. Given a trajectories of exogenous states $(S_0^{(i)}, S_1^{(i), \text{exo}}, \dots, S_{H-1}^{(i), \text{exo}})$, we can rollout a synthetic trajectory under π and \hat{P}^{end} : $\tau_D^{(i)} = (S_0^{(i)}, A_0^{(i)}, S_1^{(i)}, A_1^{(i)}, \dots)$ where $A_t^{(i)} \sim \pi(S_t^{(i)})$, $S_{t+1}^{(i), \text{end}} \sim \hat{P}^{\text{end}}(S_t^{(i)}, A_t^{(i)})$ and $S_{t+1}^{(i)} = [S_{t+1}^{(i), \text{exo}}, S_{t+1}^{(i), \text{end}}]$. For $R(\tau_D^{(i)}) := \sum_{t=0}^{H-1} r(S_t^{(i)}, A_t^{(i)})$, the average return over the N synthetic trajectories $\hat{J}(\pi, M) = \frac{1}{N} \sum_{i=1}^N R(\tau_D^{(i)})$ is an estimator of $J(\pi, M)$. This method is simple, but very useful because we can do hyperparameter selection with only the offline dataset without introducing extra hyperparameters.

We can bound the policy evaluation error by Hoeffding’s inequality. More sophisticated bounds for policy evaluation can be found in (Thomas et al., 2015).

Theorem 2. Under Assumption 1, given a deterministic policy π , we have that with probability at least $1 - \zeta$

$$\left| \hat{J}(\pi, M) - J(\pi, M) \right| \leq v_{max} \left(H \varepsilon_{air} + H \varepsilon_p + \sqrt{\frac{\ln(2/\zeta)}{2N}} \right).$$

The results suggests that if we have a sufficient number of trajectories and small ε_{air} and ε_p , then $\hat{J}(\pi, M)$ is a good estimator for $J(\pi, M)$. Even though the estimator is biased and not consistent, we find it provides sufficient information for hyperparameter selection in our experiments.

Theorem 2 only holds for a given policy that is independent of the offline data. If we want to evaluate the output policy, which depends on the data, we need to apply an union bound for all deterministic policies. In that case, the sampling error term becomes $\tilde{O}(\sqrt{|\mathcal{S}|/N})$. To avoid the dependence on the state space, an alternative is to split the data

into two subsets: one subset is used to obtain an output policy and another subset is used to evaluate the output policy.

7. Simulation Experiments

We evaluate the performance of the FQI-based algorithm in two simulated environments with the AIR property: an optimal order execution problem and an inventory management problem. We also tested the other algorithms described in Section 5.1, for completeness and to contrast to FQI-AIR. We include these results in Appendix C. FQI-AIR is notably better than these other approaches, and so we focus on it as the representative algorithm that exploits the AIR property.

The first goal of these experiments is to demonstrate that existing offline RL algorithms fail to learn a good policy for some natural data collection policies, while our proposed algorithm returns a near-optimal policy. To demonstrate this, we test three data collection policies: (1) a random policy which is designed to give a small concentration coefficient, (2) a learned near-optimal policy obtained using DQN with online interaction, which covers an optimal policy reasonably well, and (3) a constant policy which, in theory, has an infinite concentration coefficient due to missing state-action pairs and does not cover an optimal policy. The second goal is to validate the policy evaluation analysis with a varying number of trajectories N and ε_{air} .

7.1 Environments

We investigated the behavior of the algorithms on two simulated environments that mimic two real-world problems that satisfy the AIR property: optimal order execution and inventory management. In the **optimal order execution problem**, the task is to sell $M = 10$ shares within $H = 100$ steps. The stock prices X_1, \dots, X_h are generated by an ARMA(2, 2) process and scaled to the interval $[0, 1]$. Specifically, the ARMA(2, 2) process is

$$X_t = \tilde{X}_t/20 + 1/2 \quad \text{where} \quad \tilde{X}_t = c + \varepsilon_t + \sum_{i=1}^2 \varphi_i \tilde{X}_{t-i} + \sum_{i=1}^2 \theta_i \varepsilon_{t-i},$$

$\varphi_1 \sim \mathcal{U}(-0.9, 0.0)$, $\varphi_2 \sim \mathcal{U}(0.0, 0.9)$ and $\theta_i \sim \mathcal{U}(-0.5, 0.5)$ for $i = 1, 2$. The scaling parameters are chosen so that the process is stable and the price is in the interval $[0, 1]$. The endogenous variable P_h is the number of shares left to sell. To construct state, we use the most recent $K = 3$ prices with the most recent endogenous variable, that is, $S_h = (X_{h-2}, X_{h-1}, X_h, P_h)$. The action space is $\mathcal{A} = [5]$. The reward function is the stock price X_h multiplied by the number of selling shares $\min\{A_h, P_h\}$.

We consider both a setting with $\varepsilon_{air} = 0$ and $\varepsilon_{air} > 0$, as well as different data generating policies. When the number of selling shares is greater than 0, the stock price drops by 10% with probability ε_{air} . For $\varepsilon_{air} = 0$, this means selling shares has no impact on the stock price. When $\varepsilon_{air} > 0$, it does, allowing us to test how robust FQI-AIR is to some violation of the AIR property. The random policy used in the environment chooses 0 with probability 75% and choose $1, \dots, 5$ with probability 5%. The constant policy always chooses action 0.

We design an **inventory management problem** based on existing literature (Kunnumkal & Topaloglu, 2008; Van Roy, Bertsekas, Lee, & Tsitsiklis, 1997). The task is to

control the inventory of a product over $H = 100$ stages. At each stage, we observe the inventory level X_t (endogenous) and the previous demand D_{t-1} (exogenous) and choose to place a order $A_t \in [10]$. The inventory level evolves as: $X_{t+1} = (X_t + A_t - D_t)^+$. The reward function is $cA_t + h(X_t + A_t - D_t)^+ - b(-X_t - A_t + D_t)^+$ where c is the order cost, h is the holding cost, and b is the cost for lost sale. We use $c = 0.1$, $h = 0.25$ and $b = 1.0$ in the experiment. To make sure the reward is bounded, we clip the reward at a large negative number -100 . The endogenous variable is the inventory level, which can be as large as 1000, so we restrict FQI-AIR to sweep only for a subset of the endogenous space, that is, $[15] \subset \mathcal{S}^{\text{end}}$.

As before, we consider both $\varepsilon_{air} = 0$ and $\varepsilon_{air} > 0$, which in this case impacts the demand. The demand $D_t = (\tilde{D}_t)^+$ where \tilde{D}_t follows a normal distribution with mean μ and $\sigma = \mu/3$ and $(d)^+ := \max\{d, 0\}$. In the beginning of each episode, μ is sampled from a uniform distribution in the interval $[3, 9]$. When the order is greater than 0, the mean of the demand distribution decreases or increases by 10% with probability $\varepsilon_{air}/2$ respectively.

Again, we consider three different data generating policies. The random policy used in the environment is a policy which chooses a value $\tilde{A}_t \in [D_{t-1} - 3, D_{t-1} + 3]$ uniformly and then choose the action $A_t = \max\{\min\{\tilde{A}_t, 10\}, 0\}$. The constant policy always chooses action $A_t = \min\{D_{t-1}, 10\}$. The near-optimal policy is obtained using DQN with online interaction, for both environments.

There is an important nuance for the inventory problem. In this problem, the endogenous transition and reward depends on the next exogenous variable. Fortunately, we can generalize the definition of exogenous MDPs such that the endogenous transition is $P^{\text{end}} : \mathcal{S} \times \mathcal{A} \times \mathcal{S}^{\text{exo}} \rightarrow \Delta(\mathcal{S}^{\text{end}})$ and the reward is $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S}^{\text{exo}} \rightarrow [0, r_{max}]$. We assume we have an approximate endogenous model $\hat{P}^{\text{end}} : \mathcal{S} \times \mathcal{A} \times \mathcal{S}^{\text{exo}} \rightarrow \Delta(\mathcal{S}^{\text{end}})$ such that $D_{TV}(P^{\text{end}}(s, a, s^{\text{exo}}), \hat{P}^{\text{end}}(s, a, s^{\text{exo}})) \leq \varepsilon_p$ for any $(s, a, s^{\text{exo}}) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}^{\text{exo}}$. With these changes, the algorithms and the theoretical analysis naturally extend to the new definition of exogenous MDPs.

7.2 Algorithm Details

We compare FQI-AIR to FQI, MBS-QI (Liu et al., 2020), CQL (Kumar et al., 2020), IQL (Kostrikov et al., 2022). As we discussed in the previous sections, FQI is expected to work well when the concentration coefficient is small. MBS-QI is expected to perform well when the data covers an optimal policy. CQL and IQL are strong baselines which have been shown to be effective empirically for discrete-action environments such as Atari games.

We had several choices to make for the baseline algorithms. MBS-QI requires density estimation for the data distribution μ . For the optimal order execution problem, we use state discretization and empirical counts to estimate the data distribution as used in the original paper. For the inventory problem, the state space is already discrete so there is no need for discretization. We show the results with the best threshold b from the set $\{0.002, 0.001, 0.0001, 0.00005\}$. Note that it is possible that there is no data for some states (or state discretization) visited by the output policy, and for these states, all action values are all zero. To break ties, we allow MBS-QI to choose an action uniformly at random. For CQL, we add the $\text{CQL}(\mathcal{H})$ loss with a weighting α when updating the action values. We show the results with the best α from the set $\{0.1, 0.5, 1.0, 5.0\}$ as suggested in the original




Figure 2: Comparison between algorithms in the optimal order execution problem and inventory management problem, for $\varepsilon_{air} = 0$. The gray lines show the performance of the data collection policies. Results are averaged over 30 runs with error bars for one standard error.

paper. For IQL, we show the results with the best τ from the set $\{0.7, 0.8, 0.9\}$ and β from the set $\{10.0, 3.0, 1.0\}$.

We use the same value function approximation for all algorithms in our experiments: two-layers neural networks with hidden size 128. The neural networks are optimized by Adam (Kingma & Ba, 2014) or RMSprop with learning rate in the set $\{0.001, 0.0003, 0.0001\}$. All algorithms are trained for 100 iterations. We also tried training the comparator algorithms for longer, but it did not improve their performance.

The hyperparameters for FQI-AIR are selected based on $\hat{J}(\pi, M)$, which only depends on the dataset. The hyperparameters for comparator algorithms are selected based on $J(\pi, M)$ —which should be a large advantage—estimated by running the policy π on the true environment M with 100 rollouts.

7.3 Policy Performance When $\varepsilon_{air} = 0$

Figure 2 shows the performance of our algorithm and the comparator algorithms with a different number of trajectories $N = \{1, 5, 10, 25, 50, 100, 200\}$ and $\varepsilon_{air} = 0$. Our algorithm outperforms other algorithms for all data collection policies. This result is not too surprising,

as FQI-AIR is the only algorithm to exploit this important regularity in the environment; but nonetheless it shows how useful it is to exploit this AIR property when it is possible.

We can first look more closely at the optimal order execution results. MBS performs slightly better than FQI, however, we found it performs better because the tie-breaking is done with a uniform random policy especially under the constant policy dataset.² CQL and IQL fails when the data collection policy is far from optimal (constant policy) and perform reasonably with a learned policy. FQI-AIR exploits the AIR property, and so is robust to different data collection policies. The results show that exploiting the AIR property is critical for the robust performance.

We see similar patterns for the inventory management problem. FQI-AIR outperforms the other algorithms for all data collection policies. CQL and IQL perform well in this environment. MBS outperforms FQI under the learned policy, but FQI outperforms MBS under the random policy. The results match the expectation that FQI performs well with an uniform data and MBS-QI performs well with an expert data.

7.4 Policy Performance with a Large Value of ε_{air}

Now we consider the impact of using these algorithms when $\varepsilon_{air} > 0$. We should expect FQI-AIR to be most impacted, as the other algorithms do not exploit the AIR property. We vary ε_{air} from 0.1 to 0.8 and find that the results are similar to those with small ε_{air} . FQI-AIR still significantly outperforms other offline methods.

Figure 3 shows the result with $\varepsilon_{air} = 0.8$ where the performance of all algorithms drop significantly. In theory, FQI-AIR can have a large performance loss with large ε_{air} , however, FQI-AIR still consistently outperforms other baselines in our experiments, except for the inventory management problem with the learned policy. This is because the divergence between the true exogenous transition and the synthetic exogenous transition in FQI-AIR does not occur at every time step even when ε_{air} is large. For example, in the optimal order execution problem, the divergence can only happen when we sell a positive number of shares. The theoretical result is the worst-case analysis where the divergence can occur at every time step and we suffer r_{max} loss every time the divergence occurs. Therefore, the experiment results suggest that these practical problems considered in the paper are not the worst case and FQI-AIR can perform well even with large ε_{air} .

7.5 Results for Policy Evaluation

To validate the policy evaluation analysis, we investigate the difference $|\hat{J}(\pi, M) - J(\pi, M)|$ with $\varepsilon_{air} \in \{0, 0.05, 0.1, 0.2, 0.4\}$ and $N = \{1, 5, 25, 100, 200\}$ where π is the output policy of FQI-AIR. We show the 90th percentile of the difference for each combination of ε_{air} and N over 90 data points (30 runs under each data collection policy) in Figure 4. The 90th percentiles scale approximately linearly with ε_{air} and inversely proportional to N . The results suggest that the dependence on ε_{air} is linear and the policy evaluation error goes to zero at a sublinear rate, which reflects the bound of Theorem 2.

2. A uniform policy in this environment can achieve a performance $J(\pi, M) \approx 5$.

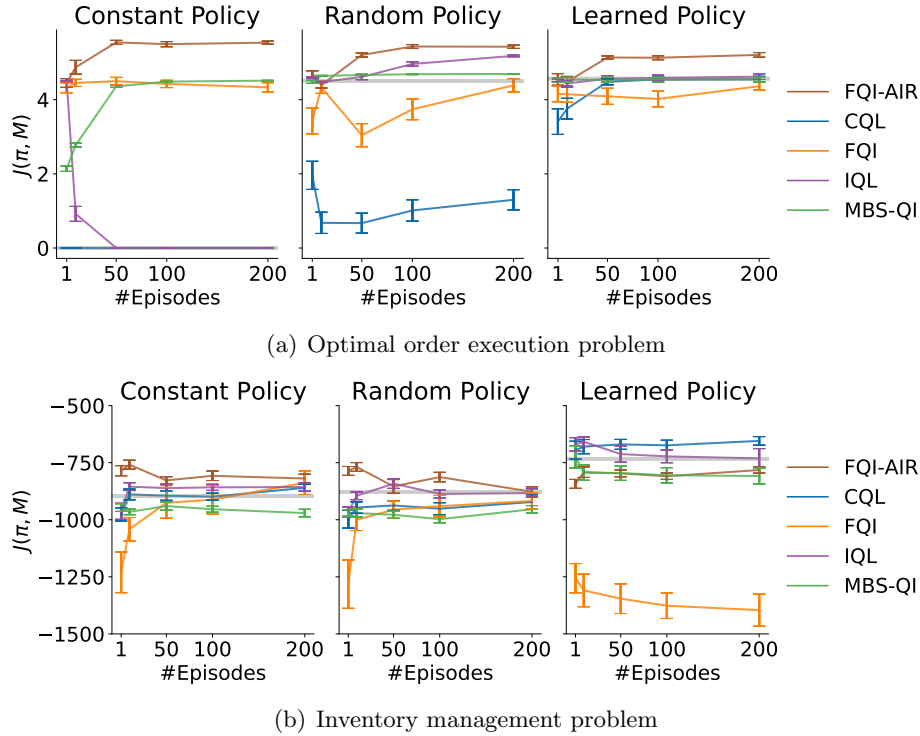


Figure 3: Comparison between algorithms in two simulation problems with $\varepsilon_{air} = 0.8$. The gray lines show the performance of the data collection policies. The results are averaged over 30 runs with error bars representing one standard error.

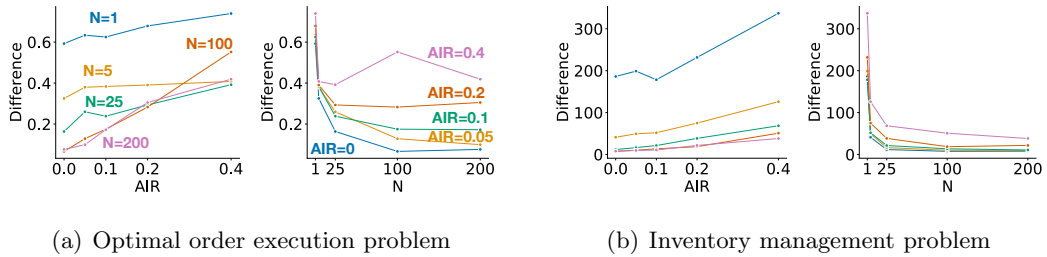


Figure 4: The 90th percentile, over 90 runs, of the difference $|\hat{J}(\pi, M) - J(\pi, M)|$, with varying N and ε_{air} . The difference should be lower for a larger N and smaller ε_{air} .

8. Real World Experiments

To demonstrate the practicality of the proposed algorithm, we evaluate the proposed algorithm for two real world experiments: (1) *Bitcoin*: an optimal order execution for the bitcoin market, and (2) *Prius*: a hybrid car control problem. For the Bitcoin experiment,

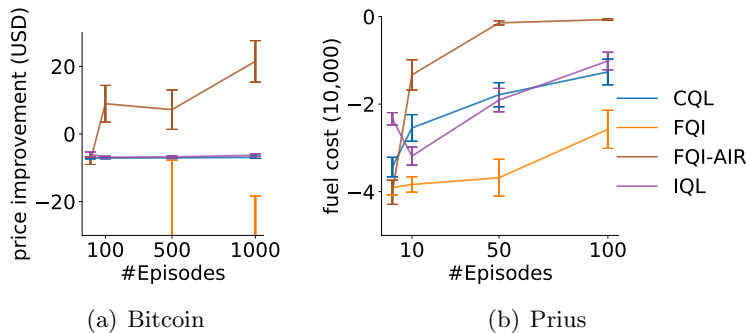


Figure 5: Performance on real world datasets. For (a), the numbers represent the total selling price minus the average price. For (b), the numbers represent the fuel cost with a penalty for not maintaining a desired batter level. Results are averaged over 30 runs, shown with one standard error.

we use historical prices of bitcoin.³ The problem is to sell one bitcoin within 60 steps where each step corresponds to 10 minutes in real world. On each step, the agent chooses to sell some numbers of bitcoins in $\{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$. Each episode corresponds to 10 hours, with a start state chosen from a random time step in the data (consisting of 300 days). The exogenous state contains the most recent 60 closing prices, and the endogenous state contains the number of shares left to be sold. We collect an offline dataset by running a trained policy by DQN for N episodes, and report performance of the output policy for the testing period (about 41 days).

For the Prius experiment, we use the hybrid car environment from Lian et al. (2020).⁴ The agent can switch between using fuel or battery, with the goal to minimize fuel consumption while maintaining a desired battery level. The exogenous state is the driving patterns and the endogenous state contains the state of charge and the previous action. We collect the offline dataset by running a learned policy with 10 different driving patterns, and test on 12 driving patterns. To better mimic the real-world, where we would not have a random policy or constant policy, we use the learned policy from DQN as the data collection policy. Further, now that the state space is larger, we run FQI-AIR where we randomly sample endogenous states and actions, rather than sweeping through all endogenous states and actions.

FQI-AIR performs significantly better than CQL, IQL and FQI, as shown in Figure 5. MBS-QI does not scale to high-dimensional continuous state spaces, and so is excluded. These results highlight that FQI-AIR can scale to high-dimensional continuous state space and large endogenous state space.

3. The bitcoin data is downloaded from the kaggle competition <https://www.kaggle.com/c/g-research-crypto-forecasting/data>.

4. We used their code at https://github.com/lryz0612/DRL-Energy-Management/blob/master/Prius_model_new.py.

9. Learning an Endogenous Model in AIR-MDPs

In the previous experiments, we assume we are given the endogenous models. In this section, we investigate the impact of using an approximate endogenous model learned from offline data. We perform this experiment in the hybrid car environment, which reflects a setting where the endogenous dynamics might in fact not be known and would be useful to estimate from data. We use a neural network to approximate the endogenous state and reward model, and run FQI-AIR with the learned endogenous model.

Let us first reason about when it might be feasible to learn a reasonably accurate endogenous model. In the worse case, learning an accurate endogenous and reward model would require data coverage for the entire state space. However, in many practical scenarios, the endogenous model can be easy to learn and does not require full data coverage. For example, in the optimal order execution problem, the endogenous dynamics does not depend on the exogenous variables, as a result, we only need coverage for the endogenous state. In the Prius environment, collecting data from just one driving cycle is sufficient to learn a good endogenous model, as we will demonstrate in these experiments.

We first collect a dataset from the hybrid car environment by running a random policy in one of the driving cycles and a deterministic policy for the other driving cycles. This data generation approach mimics a scenario we might see in practice. In the factory, we might have a test system for which it is acceptable to try many different actions (using gas or the battery), and so get a more varied dataset for learning the endogenous model. We would only get this data from one limited course (one driving cycle). The rest of the data would be collected in the wild, where the deployed solution should not be exploring many actions and should largely be deterministic.

We also test two model-based baselines (MB): (1) The first baseline has full knowledge of the reward and endogenous models, and learns the exogenous model from offline data without exploiting the AIR property. The algorithm is similar to Algorithm 1 but s_{h+1}^{exo} is generated from the learned endogenous model. (2) The second baseline does not have knowledge of the reward, endogenous models and exogenous models. It learns a full model to from a state-action pair to the next state. For these model-based baselines, the model is parameterized by a two-layer neural network and learned by minimizing the ℓ_2 distance between predicted states and next states recorded in the data. The transitions in these environments are deterministic, so it is appropriate to learn an expectation model.

In Figure 6 (a), we perform an ablation study to compare FQI-AIR and MB with the true endogenous model or a learned endogenous model. The result shows that (1) MB with the true endogenous model performs slightly worse than FQI-AIR with a small data size. (2) FQI-AIR with a learned endogenous model perform worse than FQI-AIR, however, it outperforms IQL and MB without the true endogenous model. (3) MB with a learned endogenous model performs worse than FQI-AIR with a learned endogenous model. This suggests that it is useful to separate the exogenous state and endogenous state especially when we need to learn an endogenous model.

Next, we test FQI-AIR when learning the endogenous model only from a more limited dataset: a dataset based solely on one cycle. We collect a dataset from the hybrid car environment by running a random policy in one of the driving cycles for 500 episodes. This reflects a practical scenario that we can just run our vehicle in a closed area and still are

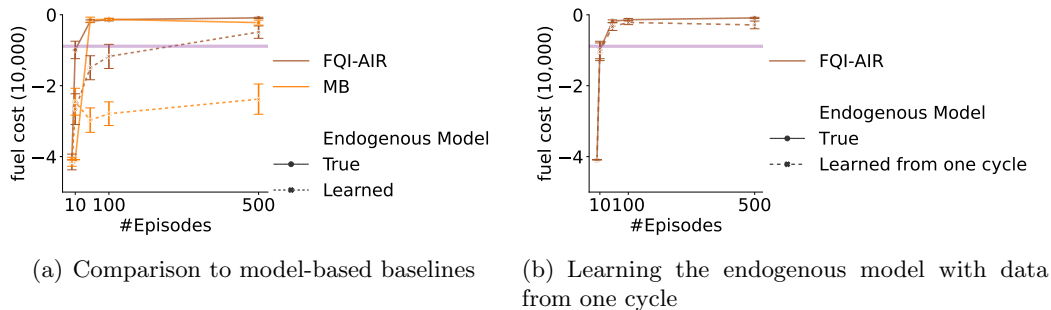


Figure 6: Experiment on the Prius dataset. We include IQL with 500 episodes as a baseline (purple line). Results are averaged over 30 runs, shown with one standard error.

able to obtain a good endogenous model for running FQI-AIR. Figure 6 (b) shows that FQI-AIR with the learned endogenous model performs well and is close to FQI-AIR with true endogenous model.

10. Conclusion

In this paper, we aim to understand whether and when offline RL is feasible for real world problems. It is known that, without extra assumptions on MDPs, offline RL requires exponential numbers of samples to obtain a nearly optimal policy with high probability even if we have a good data distribution. As a result, we must make assumptions on MDPs to make learning feasible (learning with polynomial sample complexity). However, common assumptions can be impractical as discussed in Section 3. Therefore, our goal in this paper is to study an MDP property that (1) is realistic for several important real world problems and (2) makes offline RL feasible.

We introduced an MDP property, which we call Action Impact Regularity (AIR). We developed an algorithm for MDPs satisfying AIR, that (1) has strong theoretical guarantees on the supoptimality, without making assumptions about concentration coefficients or data coverage, (2) provides a simple way to select hyperparameters offline, without introducing extra hyperparameters and (3) is simple to implement and computationally efficient. We showed empirically that the proposed algorithm significantly outperforms existing offline RL algorithms, across two simulated environments and two real world environments.

Appendix A. Experiment Details

In this section, we provide more experimental details.

Algorithm Details The original MBS-QI algorithm does not work with negative rewards. The reward in the inventory problem is in the range $[-100, 0]$, so we modify the pessimistic value for MBS-QI for the problem:

$$\tilde{q}_h(s, a) := \begin{cases} q_h(s, a), & \text{if } \hat{\mu}(s, a) \geq b \\ -10000, & \text{if } \hat{\mu}(s, a) < b \end{cases}$$

instead of $\tilde{q}_h(s, a) := \mathbb{I}\{\hat{\mu}(s, a) \geq b\}q_h(s, a)$ where $\hat{\mu}$ is the estimated data distribution.

Data Collection Policies To collect data for the learned policy, we first train a DQN algorithm for 1000 episodes with online interaction with the underlying environment. The DQN parameters (that is, learning rate and optimizer) are chosen based on estimated $J(\pi, M)$. After training, we collect data by running the trained policy.

Appendix B. Theoretical Analysis

We first provide definitions and lemmas that would be useful for proving our main theorems. Given $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$, β_h be a probability measure on $\mathcal{S} \times \mathcal{A}$ and $p > 0$, define $\|q\|_{p, \beta_h}^p = \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \beta_h(s, a) |q(s, a)|^p$. Given $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$, we define the Bellman evaluation operator for a given policy π :

$$(\mathcal{T}_\pi q)(s, a) := r(s, a) + \sum_{s' \in \mathcal{S}} P(s, a, s') \sum_{a' \in \mathcal{A}} \pi(a' | s') q(s', a').$$

Lemma 1. Given a MDP M , suppose the sequence of value function (q_0, \dots, q_{H-1}) satisfies that $\|\mathcal{T}q_{h+1} - q_h\|_{2, d_h^\pi} \leq \varepsilon$ for all policy π with $q_H = 0$. Let π be the greedy policy with respect to (q_0, \dots, q_{H-1}) , then we have that

$$J(\pi, M) \geq J(\pi_M^*, M) - (H + 1)H\varepsilon.$$

Proof. By the performance difference lemma (Lemma 5.2.1 of Kakade (2003) or Chen and Jiang (2019)), we get

$$\begin{aligned} & J(\pi^*, M_D) - J(\pi, M_D) \\ &= \mathbb{E}^{\pi^*} \left[\sum_{h=0}^{H-1} q^*(S_h, A_h) - q^*(S_h, \pi_h(S_h)) \right] \\ &\leq \mathbb{E}^{\pi^*} \left[\sum_{h=0}^{H-1} q^*(S_h, A_h) - q_h(S_h, A_h) + q_h(S_h, \pi_h(S_h)) - q^*(S_h, \pi_h(S_h)) \right] \\ &\leq \mathbb{E}^{\pi^*} \left[\sum_{h=0}^{H-1} |q^*(S_h, A_h) - q_h(S_h, A_h)| + |q_h(S_h, \pi_h(S_h)) - q^*(S_h, \pi_h(S_h))| \right] \\ &\leq \sum_{h=0}^{H-1} \|q^* - q_h\|_{1, d_h^{\pi^*} \pi^*} + \|q^* - q_h\|_{1, d_h^{\pi^*} \pi_h} \\ &\leq \sum_{h=0}^{H-1} \|q^* - q_h\|_{2, d_h^{\pi^*} \pi^*} + \|q^* - q_h\|_{2, d_h^{\pi^*} \pi_h} \end{aligned}$$

where d_h^π is the state-action distribution at horizon h induced by policy π . The first inequality follows because π_h is greedy with respect to q_h . We consider a state-action distribution β_0 that is induced by some policy, then

$$\begin{aligned} \|q^* - q_0\|_{2, \beta_0} &\leq \|\mathcal{T}q^* - \mathcal{T}q_1 + \mathcal{T}q_1 - q_0\|_{2, \beta_0} \\ &\leq \|\mathcal{T}q^* - \mathcal{T}q_1\|_{2, \beta_0} + \|\mathcal{T}q_1 - q_0\|_{2, \beta_0} \\ &\leq \|q^* - q_1\|_{2, \beta_1} + \varepsilon \end{aligned}$$

where $\beta_1(s', a') = \sum_{s,a} \beta_0(s, a) P(s, a, s') \mathbb{I}\{a' = \arg \max_{a'' \in \mathcal{A}} (q^*(s', a'') - q_1(s', a''))^2\}$ is also induced by some policy. The first inequality follows by the fact that q^* is the fixed point of the operator \mathcal{T} . We can recursively apply the same process for $\|q^* - q_h\|_{2, \beta_h}$, $h > 0$, and we can get

$$\|q^* - q_h\|_{2, \beta_h} \leq (H - h)\varepsilon.$$

Plug in the inequality to the performance difference lemma, we get

$$J(\pi^*, M_D) - J(\pi, M_D) \leq \sum_{h=0}^{H-1} 2(H - h)\varepsilon = (H + 1)H\varepsilon.$$

□

The simulation lemma was first introduced for discounted setting in (Kearns & Singh, 2002), and here we prove a modified version of the simulation lemma for the finite horizon MDPs.

Lemma 2 (Finite Horizon Simulation Lemma). Given a policy π , let $M = (\mathcal{S}, \mathcal{A}, P, r, H, \mu)$ and $\hat{M} = (\mathcal{S}, \mathcal{A}, \hat{P}, r, H, \mu)$ be two finite horizon MDPs and π be a policy satisfying that for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$, $\|P(s, a) - \hat{P}(s, a)\|_1 \leq \varepsilon$. Then, $|v_M^\pi(s) - v_{\hat{M}}^\pi(s)| \leq \varepsilon \frac{H^2 r_{max}}{2}$ for $s \in \mathcal{S}_0$.

Proof. For all $s \in \mathcal{S}_0, a \in \mathcal{A}$,

$$\begin{aligned} & |v_M^\pi(s) - v_{\hat{M}}^\pi(s)| \\ &= |r_\pi(s) + \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} P(s, a, s') v_M^\pi(s') - r_\pi(s) - \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} \hat{P}(s, a, s') v_M^\pi(s')| \\ &= | \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} P(s, a, s') v_M^\pi(s') - \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} \hat{P}(s, a, s') v_M^\pi(s') | \\ &= | \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} P(s, a, s') v_M^\pi(s') - \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} \hat{P}(s, a, s') v_M^\pi(s') \\ &\quad + \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} \hat{P}(s, a, s') v_M^\pi(s') - \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}_1} \hat{P}(s, a, s') v_{\hat{M}}^\pi(s') | \\ &\leq \sum_{a \in \mathcal{A}} \pi(a|s) | \sum_{s' \in \mathcal{S}_1} (P(s, a, s') - \hat{P}(s, a, s')) v_M^\pi(s') | + \max_{s' \in \mathcal{S}_1} |v_M^\pi(s') - v_{\hat{M}}^\pi(s')| \\ &\leq \varepsilon(H - 1)r_{max} + \max_{s' \in \mathcal{S}_1} |v_M^\pi(s') - v_{\hat{M}}^\pi(s')| \\ &\leq \varepsilon(H - 1)r_{max} + \varepsilon_P(H - 2)r_{max} + \cdots + \varepsilon_P 1 r_{max} \\ &\leq \varepsilon \frac{H^2 r_{max}}{2} = \varepsilon \frac{H v_{max}}{2} \end{aligned}$$

The first equality follows from the Bellman equation. The second and third inequalities follow from that $v_M^\pi(s)$ is at most $(H - h)r_{max}$ for $s \in \mathcal{S}_h$. □

Lemma 3. Let $M = (\mathcal{S}, \mathcal{A}, P^{exo}, P^{end}, r, H)$ be an ε_{air} -AIR MDP and $M_b = (\mathcal{S}, \mathcal{A}, \hat{P}^{exo}, \hat{P}^{end}, r, H)$ with $D_{TV}(P^{end}(s, a), \hat{P}^{end}(s, a)) \leq \varepsilon_p$, then for any policy π ,

$$|J(\pi, M) - J(\pi, M_b)| \leq v_{max} H (\varepsilon_{air} + \varepsilon_p).$$

Proof. Since M is ε_{air} -AIR, we have

$$D_{TV} (P^{\text{exo}}(s^{\text{exo}}, a), P^{\text{exo}}(s^{\text{exo}}, a')) = \frac{1}{2} \|P^{\text{exo}}(s^{\text{exo}}, a) - P^{\text{exo}}(s^{\text{exo}}, a')\|_1 \leq \varepsilon_{air}$$

Let $e(s, a, s'^{\text{end}}) = P^{\text{end}}(s, a, s'^{\text{end}}) - \tilde{P}^{\text{end}}(s, a, s'^{\text{end}})$, we know $D_{TV} (P^{\text{end}}(s, a), \tilde{P}^{\text{end}}(s, a)) = \frac{1}{2} \|e(s, a, \cdot)\|_1 = \frac{1}{2} \sum_{s'^{\text{end}}} |e(s, a, s'^{\text{end}})| \leq \varepsilon_p$. For any $s = (s^{\text{exo}}, s^{\text{end}}) \in \mathcal{S}^{\text{exo}} \times \mathcal{S}^{\text{end}}$ and $a \in \mathcal{A}$, we have

$$\begin{aligned} & \|P(s, a) - \tilde{P}(s, a)\|_1 \\ &= \sum_{s'} |P([s^{\text{exo}}, s^{\text{end}}], a, s') - \tilde{P}([s^{\text{exo}}, s^{\text{end}}], a, s')| \\ &= \sum_{s'} |P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) P^{\text{end}}(s, a, s'^{\text{end}}) - \tilde{P}^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) \tilde{P}^{\text{end}}(s, a, s'^{\text{end}})| \\ &= \sum_{s'} |P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) (\tilde{P}^{\text{end}}(s, a, s'^{\text{end}}) + e(s, a, s'^{\text{end}})) - \tilde{P}^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) \tilde{P}^{\text{end}}(s, a, s'^{\text{end}})| \\ &= \sum_{s'} |\tilde{P}^{\text{end}}(s, a, s'^{\text{end}}) [P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) - \tilde{P}^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}})] + P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) e(s, a, s'^{\text{end}})| \\ &\leq \sum_{s'} \tilde{P}^{\text{end}}(s, a, s'^{\text{end}}) |P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) - \sum_{a'} \pi'(a' | s^{\text{exo}}) P^{\text{exo}}(s^{\text{exo}}, a', s'^{\text{exo}})| + \\ &\quad \sum_{s'} P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) |e(s, a, s'^{\text{end}})| \\ &\leq \sum_{a'} \pi'(a | s^{\text{exo}}) \sum_{s'^{\text{exo}}} |P^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) - P^{\text{exo}}(s^{\text{exo}}, a', s'^{\text{exo}})| + \sum_{s'^{\text{end}}} |e(s, a, s'^{\text{end}})| \\ &\leq 2\varepsilon_{air} + 2\varepsilon_p. \end{aligned}$$

The first inequality follows by writing $\tilde{P}^{\text{exo}}(s^{\text{exo}}, a, s'^{\text{exo}}) = \sum_{a' \in \mathcal{A}} \pi'(a' | s^{\text{exo}}) P^{\text{exo}}(s^{\text{exo}}, a', s'^{\text{exo}})$ where $\pi'(a' | s^{\text{exo}}) = \frac{\mathbb{P}_M^{\pi_b}(S_h^{\text{exo}} = s_h^{\text{exo}}, A_h = a')}{\mathbb{P}_M^{\pi_b}(S_h^{\text{exo}} = s_h^{\text{exo}})}$ when the denominator is nonzero, and otherwise let $\pi'(\cdot | s_h^{\text{exo}})$ be an arbitrary distribution. Applying Lemma 2, we get

$$|J(\pi, M) - J(\pi, M_b)| = \sum_{s_0} \mu(s_0) |v_M^\pi(s_0) - v_{M_b}^\pi(s_0)| \leq v_{max}(\varepsilon_{air} + \varepsilon_p)H.$$

□

Theorem 1. Under Assumption 1 and 2, let π_M^* be an optimal policy in M and π the output policy of FQI-AIR, then with probability at least $1 - \zeta$

$$\begin{aligned} J(\pi, M) &\geq J(\pi_M^*, M) - 2v_{max}H(\varepsilon_{air} + \varepsilon_p) - \\ &(H + 1)H \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{72v_{max}^2 \ln(H|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{apx} \right). \end{aligned}$$

Proof. First fix a horizon $h \in [H]$ and $f' = q_{h+1} \in \mathcal{F}$. Define

$$\begin{aligned} R(f) &= \|\mathcal{T}f' - f\|_{2, \tilde{v}_h}^2 \\ &= \sum_{s_h^{\text{exo}} \in \mathcal{S}_h^{\text{exo}}} \sum_{s_h^{\text{end}} \in \mathcal{S}_h^{\text{end}}} \sum_{a \in \mathcal{A}} \frac{\tilde{v}_h(s_h^{\text{exo}})}{|\mathcal{S}^{\text{end}}| |\mathcal{A}|} (f(s_h^{\text{exo}}, s_h^{\text{end}}, a) - r - \mathbb{E}[\max_{a'} f'(S_{h+1}^{\text{exo}}, S_{h+1}^{\text{end}}, a')])^2 \end{aligned}$$

and

$$R_n(f) = \frac{1}{N} \sum_{i=1}^N \sum_{s_h^{\text{end}} \in \mathcal{S}_h^{\text{end}}} \sum_{a \in \mathcal{A}} \frac{1}{|\mathcal{S}^{\text{end}}| |\mathcal{A}|} (f(s_h^{(i), \text{exo}}, s_h^{\text{end}}, a) - r - \max_{a'} f'(s_{h+1}^{(i), \text{exo}}, s_{h+1}^{\text{end}}, a'))^2$$

where $s_{h+1}^{\text{end}} \sim \hat{P}^{\text{end}}(s_h^{(i), \text{exo}}, s_h^{\text{end}}, a)$. Let

$$\hat{f} = \arg \min_{f \in \mathcal{F}} R_n(f), \quad \tilde{f} = \arg \min_{f \in \mathcal{F}} R(f),$$

our goal is to bound $R(\hat{f})$ with high probability. This is similar to bounding the generalization error in the statistical learning literature. We follow the proof technique of Lemma A.11 in Agarwal, Jiang, Kakade, and Sun (2019) to bound the excess risk.

Fix $u = (s_h^{\text{end}}, a) \in \mathcal{U} := \mathcal{S}^{\text{end}} \times \mathcal{A}$, let $x_i^u = (s_h^{\text{exo}}, s_h^{\text{end}}, a)$ and $y_i^u = r(s_h^{\text{exo}}, s_h^{\text{end}}, a) + \max_{a'} f'(s_{h+1}^{\text{exo}}, s_{h+1}^{\text{end}}, a')$ with $(x_i^u, y_i^u) \sim \nu$, $f^*(x_i^u) = \mathbb{E}[y_i^u | x_i^u]$. Note that our goal is to bound $R(\hat{f}) = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{E}[(\hat{f}(x_i^u) - f^*(x_i^u))^2]$.

First note that

$$\mathbb{E}[(f(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] = \mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2]$$

and

$$\begin{aligned} \mathbb{V}^{u, f} &:= \mathbb{V}[(f(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] \\ &\leq \mathbb{E}[(f(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2]^2 \\ &\leq 4v_{\max}^2 \mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2]. \end{aligned}$$

We can bound the deviation from the mean for all $f \in \mathcal{F}$ with one-sided Bernstein's inequality: the following holds with probability $1 - \zeta$,

$$\begin{aligned} &\mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2] - \frac{1}{N} \sum_{i=1}^N [(f(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] \\ &\leq \sqrt{\frac{2\mathbb{V}^{u, f} \ln(|\mathcal{F}|/\zeta)}{N}} + \frac{4v_{\max}^2 \ln(|\mathcal{F}|/\zeta)}{3N} \\ &\leq \sqrt{\frac{8v_{\max}^2 \mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2] \ln(|\mathcal{F}|/\zeta)}{N}} + \frac{4v_{\max}^2 \ln(|\mathcal{F}|/\zeta)}{3N}. \end{aligned}$$

We need this holds for all $u \in \mathcal{U}$. By the union bound, we have for all $u \in \mathcal{U}$

$$\begin{aligned} &\mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2] - \frac{1}{N} \sum_{i=1}^N [(f(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] \\ &\leq \sqrt{\frac{8v_{\max}^2 \mathbb{E}[(f(x_i^u) - f^*(x_i^u))^2] \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{N}} + \frac{4v_{\max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{3N}. \end{aligned}$$

Since \hat{f} is the empirical minimizer, we have

$$\begin{aligned} & \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{N} \sum_{i=1}^N [(\hat{f}(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] \\ & \leq \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{N} \sum_{i=1}^N [(\tilde{f}(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2]. \end{aligned}$$

Since $\tilde{f} \in \mathcal{F}$, we have

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N [(\tilde{f}(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] - \mathbb{E}[(\tilde{f}(x_i^u) - f^*(x_i^u))^2] \\ & \leq \sqrt{\frac{8v_{max}^2 \mathbb{E}[(\tilde{f}(x_i^u) - f^*(x_i^u))^2] \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{N}} + \frac{4v_{max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{3N}. \end{aligned}$$

Since $\frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{E}[(\tilde{f}(x_i^u) - f^*(x_i^u))^2] \leq \varepsilon_{apx}$, we have

$$\begin{aligned} & \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{N} \sum_{i=1}^N [(\tilde{f}(x_i^u) - y_i^u)^2 - (f^*(x_i^u) - y_i^u)^2] \\ & \leq \varepsilon_{apx} + \sqrt{\frac{8v_{max}^2 \varepsilon_{apx} \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{N}} + \frac{4v_{max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{3N} \\ & \leq \frac{3}{2} \varepsilon_{apx} + \frac{16v_{max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{3N}. \end{aligned}$$

The second inequality follows by the fact that $\sqrt{ab} \leq \frac{a+b}{2}$ for $a, b \geq 0$. Then,

$$\begin{aligned} & \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{E}[(\hat{f}(x_i^u) - f^*(x_i^u))^2] \\ & \leq \sqrt{\frac{8v_{max}^2 \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{E}[(\hat{f}(x_i) - f^*(x_i))^2] \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{N}} + \frac{20v_{max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{3N} + \frac{3}{2} \varepsilon_{apx}. \end{aligned}$$

Solving for the quadratic formula, we get

$$\frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{E}[(\hat{f}(x_i^u) - f^*(x_i^u))^2] \leq \frac{36v_{max}^2 \ln(|\mathcal{F}||\mathcal{U}|/\zeta)}{N} + 2\varepsilon_{apx}.$$

Now define $\tilde{d}_h^\pi(s^{\text{exo}}) = \mathbb{P}_{M_b}^\pi(S_h^{\text{exo}} = s_h^{\text{exo}})$ and $d_h^\pi(s^{\text{exo}}, s^{\text{end}}, a) = \mathbb{P}_{M_b}^\pi(S_h^{\text{exo}} = s_h^{\text{exo}}, S_h^{\text{end}} = s_h^{\text{end}}, A_h = a)$. By the construction of M_b , we have

$$\begin{aligned} & \forall \pi, s_h^{\text{exo}} \in S_h^{\text{exo}}, \frac{\tilde{d}_h^\pi(s^{\text{exo}})}{\tilde{v}_h(s^{\text{exo}})} \leq 1, \text{ and} \\ & \forall \pi, s_h^{\text{exo}} \in S_h^{\text{exo}}, s_h^{\text{end}} \in S_h^{\text{end}}, a \in \mathcal{A}, \frac{d_h^\pi(s^{\text{exo}}, s^{\text{end}}, a)}{\tilde{v}_h(s^{\text{exo}})/|\mathcal{S}^{\text{end}}||\mathcal{A}|} \leq |\mathcal{S}^{\text{end}}||\mathcal{A}|. \end{aligned}$$

Therefore, for any policy π , we have with probability $1 - \zeta$

$$\begin{aligned} \|\mathcal{T}q_{h+1} - q_h\|_{2, d_h^\pi} &\leq \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \|\mathcal{T}q_{h+1} - q_h\|_{2, \tilde{v}_h} \\ &\leq \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{36v_{\max}^2 \ln(|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{\text{apx}} \right). \end{aligned}$$

Note that this holds for a fixed $h \in [H]$ and $f' \in \mathcal{F}$. Use the union bound, we have that, for any policy π , $h \in [H]$ and $q_{h+1} \in \mathcal{F}$, with probability $1 - \zeta$

$$\|\mathcal{T}q_{h+1} - q_h\|_{2, d_h^\pi} \leq \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{72v_{\max}^2 \ln(H|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{\text{apx}} \right).$$

By Lemma 1, the output policy π satisfies

$$\begin{aligned} J(\pi_{M_b}^*, M_b) - J(\pi, M_b) \\ \leq (H+1)H \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{72v_{\max}^2 \ln(H|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{\text{apx}} \right) = \varepsilon_1. \end{aligned}$$

By Lemma 3, the output policy π satisfies

$$\begin{aligned} J(\pi_M^*, M) - J(\pi, M) \\ &= J(\pi_M^*, M) - J(\pi_M^*, M_b) + J(\pi_M^*, M_b) - J(\pi, M) \\ &\leq J(\pi_M^*, M) - J(\pi_M^*, M_b) + J(\pi_{M_b}^*, M_b) - J(\pi, M) \\ &\leq J(\pi_M^*, M) - J(\pi_M^*, M_b) + J(\pi, M_b) - J(\pi, M) + \varepsilon_1 \\ &\leq |J(\pi_M^*, M) - J(\pi_M^*, M_b)| + |J(\pi, M_b) - J(\pi, M)| + \varepsilon_1 \\ &\leq 2v_{\max}H(\varepsilon_{\text{air}} + \varepsilon_P) + (H+1)H \sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|} \left(\sqrt{\frac{72v_{\max}^2 \ln(H|\mathcal{F}||\mathcal{S}^{\text{end}}||\mathcal{A}|/\zeta)}{N}} + 2\varepsilon_{\text{apx}} \right). \end{aligned}$$

Note that $\sqrt{|\mathcal{S}^{\text{end}}||\mathcal{A}|}$ comes from the fact that we run FQI for all endogenous state and action uniformly. This term can be viewed as the concentration coefficient in the standard FQI analysis and is unavoidable. However, if we can adapt the weight on the endogenous state and action, we might be able to reduce the dependence. \square

Theorem 2. Under Assumption 1, given a deterministic policy π , we have that with probability at least $1 - \zeta$

$$\left| \hat{J}(\pi, M) - J(\pi, M) \right| \leq v_{\max} \left(H\varepsilon_{\text{air}} + H\varepsilon_P + \sqrt{\frac{\ln(2/\zeta)}{2N}} \right).$$

Proof. We first show that $R(\tau_D^{(i)}) := \sum_{t=0}^{H-1} r(s_t^{(i)}, a_t^{(i)})$ for $i \in [N]$ are i.i.d. samples with mean $J(\pi, M_b)$. Let \mathbb{P}_D^π be the probability measure on trajectories τ_D . Note that the randomness of τ_D comes from the generation of exogenous variables in by the interaction between π_b and M , and the generation of actions and endogenous variables by π and \hat{P}^{end} . Let $\mathbb{P}_{M_b}^\pi$ be the probability measure on the trajectories sampled by running π on M_b . It is sufficient to show that $\mathbb{P}_D^\pi = \mathbb{P}_{M_b}^\pi$ then $\mathbb{E}[R(\tau_D)] = \int R(\tau) d\mathbb{P}_D^\pi(\tau) = \int R(\tau) d\mathbb{P}_{M_b}^\pi(\tau) = J(\pi, M_b)$.

For all trajectories $\tau = (s_0, a_0, \dots, s_{H-1}, a_{H-1}) \in (\mathcal{S} \times \mathcal{A})^H$, we have

$$\begin{aligned} \mathbb{P}_D^\pi(\tau) &= \mathbb{P}_M^{\pi_b}(s_0, s_1^{\text{exo}}, \dots, s_{H-1}^{\text{exo}}) \mathbb{P}_D^\pi(a_0, s_1^{\text{end}}, \dots, s_{H-1}^{\text{end}}, a_{H-1} \mid s_0, s_1^{\text{exo}}, \dots, s_{H-1}^{\text{exo}}) \\ &= \mu(s_0) \pi(a_0 \mid s_0) \mathbb{P}_M^{\pi_b}(s_1^{\text{exo}} \mid s_0) \hat{P}^{\text{end}}(s_0, a_0, s_1^{\text{end}}) \pi(a_1 \mid s_1) \dots \\ &= \mu(s_0) \pi(a_0 \mid s_0) \tilde{P}(s_0, a_0, s_1) \pi(a_1 \mid s_1) \dots = \mathbb{P}_{M_b}^\pi(\tau). \end{aligned}$$

By Hoeffding’s inequality, with probability at least $1 - \zeta$,

$$\left| \hat{J}(\pi, M) - J(\pi, M_b) \right| = \left| \frac{1}{N} \sum_{i=1}^N R(\tau_D^{(i)}) - \mathbb{E}[R(\tau_D)] \right| \leq v_{\max} \sqrt{\frac{1}{2N} \ln \frac{2}{\zeta}}.$$

Finally by Lemma 3, the followings hold with probability at least $1 - \zeta$:

$$\begin{aligned} |\hat{J}(\pi, M) - J(\pi, M)| &\leq |\hat{J}(\pi, M) - J(\pi, M_b)| + |J(\pi, M_b) - J(\pi, M)| \\ &\leq v_{\max} H(\epsilon_{AIR} + \epsilon_P) + v_{\max} \sqrt{\frac{\ln(2/\zeta)}{2N}}. \end{aligned}$$

□

Appendix C. Comparing FQI-AIR to Online RL with Trajectory Simulation




Figure 7: Comparison to simulation-based algorithms. Results are averaged over 30 runs with the shaded region representing plus and minus one standard error.

We compare to two online RL algorithms, approximate policy iteration (API) and actor critic (AC) with ϵ -greedy exploration, with the trajectory simulator discussed in Section 5.1. We show the learning curves with random data collection policy and $N = 100$. Each iteration contains a sweep over the entire dataset and we evaluate the greedy policy with respect to the Q function after each iteration. The results show that FQI-AIR converges to a nearly optimal solution within a few iterations while the online RL algorithms require much more iterations to find a good policy. AC in the optimal order execution problem does not converge to a stable performance. The final performance of API and AC are also worse than the performance of FQI-AIR. These results show that online RL algorithms could be used for AIR MDPs. However, they are less direct and efficient, and they could find a slightly different solution with finite data.

References

- Abdullah, H. M., Gastli, A., & Ben-Brahim, L. (2021). Reinforcement learning based EV charging management systems—a review. *IEEE Access*, *9*, 41506–41531.
- Abernethy, J., & Kale, S. (2013). Adaptive market making via online learning. *Advances in Neural Information Processing Systems*.
- Agarwal, A., Jiang, N., Kakade, S. M., & Sun, W. (2019). *Reinforcement learning: Theory and algorithms*.
- Agarwal, A., Kakade, S., & Yang, L. F. (2020). Model-based reinforcement learning with a generative model is minimax optimal. In *Conference on Learning Theory*.
- Antos, A., Szepesvári, C., & Munos, R. (2006). Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. In *International Conference on Computational Learning Theory*.
- Babaioff, M., Immorlica, N., Kempe, D., & Kleinberg, R. (2007). A knapsack secretary problem with applications. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Springer.
- Bellemare, M. G., Veness, J., & Bowling, M. (2012). Investigating contingency awareness using Atari 2600 games. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Bertsimas, D., & Lo, A. W. (1998). Optimal control of execution costs. *Journal of financial markets*, *1*(1), 1–50.
- Brafman, R. I., & Tennenholtz, M. (2002). R-Max—a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, *3*(Oct), 213–231.
- Burhani, H., Ding, G. W., Hernandez-Leal, P., Prince, S., Shi, D., & Szeto, S. (2020). Aiden – reinforcement learning for order execution. <https://www.borealisai.com/en/blog/aiden-reinforcement-learning-for-order-execution/>. [Online; accessed 1-June-2022].
- Chen, J., & Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*.
- Chitnis, R., & Lozano-Pérez, T. (2020). Learning compact models for planning with exogenous processes. In *Conference on Robot Learning*.
- Dietterich, T., Trimponias, G., & Chen, Z. (2018). Discovering and removing exogenous state variables and rewards for reinforcement learning. In *International Conference on Machine Learning*.
- Duan, Y., Jin, C., & Li, Z. (2021). Risk bounds and Rademacher complexity in batch reinforcement learning. In *International Conference on Machine Learning*.
- Efroni, Y., Misra, D., Krishnamurthy, A., Agarwal, A., & Langford, J. (2022). Provable RL with exogenous distractors via multistep inverse dynamics. In *International Conference on Learning Representations*.

- Ernst, D., Geurts, P., & Wehenkel, L. (2005). Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6.
- Farahmand, A.-m., Szepesvári, C., & Munos, R. (2010). Error propagation for approximate policy and value iteration. *Advances in Neural Information Processing Systems*.
- Freeman, P. (1983). The secretary problem and its extensions: A review. *International Statistical Review/Revue Internationale de Statistique*.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., & Levine, S. (2020). D4RL: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Fujimoto, S., Meger, D., & Precup, D. (2019). Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*.
- Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, 66(3), 1375–1394.
- Jiang, N. (2018). PAC reinforcement learning with an imperfect model. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Kakade, S. M. (2003). *On the sample complexity of reinforcement learning*. Ph.D. thesis, University College London.
- Kearns, M., & Singh, S. (2002). Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2), 209–232.
- Kidambi, R., Rajeswaran, A., Netrapalli, P., & Joachims, T. (2020). MOREL: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kostrikov, I., Nair, A., & Levine, S. (2022). Offline reinforcement learning with implicit q-learning. In *International Conference on Learning Representations*.
- Kumar, A., Fu, J., Soh, M., Tucker, G., & Levine, S. (2019). Stabilizing off-policy Q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*.
- Kumar, A., Zhou, A., Tucker, G., & Levine, S. (2020). Conservative Q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*.
- Kunnumkal, S., & Topaloglu, H. (2008). Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research*, 56(3), 646–664.
- Laroche, R., Trichelair, P., & Des Combes, R. T. (2019). Safe policy improvement with baseline bootstrapping. In *International Conference on Machine Learning*.
- Lattimore, T., & Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Levine, S., Kumar, A., Tucker, G., & Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Lian, R., Peng, J., Wu, Y., Tan, H., & Zhang, H. (2020). Rule-interposing deep reinforcement learning based energy management strategy for power-split hybrid electric vehicle. *Energy*, 197, 117297.

- Liu, Y., Swaminathan, A., Agarwal, A., & Brunskill, E. (2020). Provably good batch reinforcement learning without great exploration. *Advances in Neural Information Processing Systems*.
- McGregor, S., Houtman, R., Montgomery, C., Metoyer, R., & Dietterich, T. G. (2017). Factoring exogenous state for model-free Monte Carlo. *arXiv preprint arXiv:1703.09390*.
- Munos, R. (2003). Error bounds for approximate policy iteration. In *International Conference on Machine Learning*.
- Munos, R. (2005). Error bounds for approximate value iteration. In *National Conference on Artificial Intelligence*.
- Munos, R. (2007). Performance bounds in L_p -norm for approximate value iteration. *SIAM journal on control and optimization*, 46(2), 541–561.
- Nevmyvaka, Y., Feng, Y., & Kearns, M. (2006). Reinforcement learning for optimized trade execution. In *International Conference on Machine Learning*.
- Shahamiri, M. (2008). Reinforcement learning in environments with independent delayed-sense dynamics. Master’s thesis, University of Alberta.
- Thomas, P., Theocharous, G., & Ghavamzadeh, M. (2015). High confidence policy improvement. In *International Conference on Machine Learning*.
- Van Roy, B., Bertsekas, D. P., Lee, Y., & Tsitsiklis, J. N. (1997). A neuro-dynamic programming approach to retailer inventory management. In *IEEE Conference on Decision and Control*.
- Wang, R., Foster, D. P., & Kakade, S. M. (2021). What are the statistical limits of offline RL with linear function approximation?. In *International Conference on Learning Representations*.
- Wu, Y., Tucker, G., & Nachum, O. (2019). Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*.
- Xiao, C., Lee, I., Dai, B., Schuurmans, D., & Szepesvari, C. (2022). The curse of passive data collection in batch reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*.
- Yin, M., Bai, Y., & Wang, Y.-X. (2021). Near-optimal provable uniform convergence in offline policy evaluation for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*.
- Zanette, A. (2021). Exponential lower bounds for batch reinforcement learning: Batch RL can be exponentially harder than online RL. In *International Conference on Machine Learning*.