

Prospect-theoretic Q-learning^{*}

Vivek S. Borkar, Siddharth Chandak^a

^a*Department of Electrical Engineering, Indian Institute of Technology Bombay, Powai, Mumbai-400076, India*

Abstract

We consider a prospect theoretic version of the classical Q-learning algorithm for discounted reward Markov decision processes, wherein the controller perceives a distorted and noisy future reward, modeled by a nonlinearity that accentuates gains and under-represents losses relative to a reference point. We analyze the asymptotic behavior of the scheme by analyzing its limiting differential equation and using the theory of monotone dynamical systems to infer its asymptotic behavior. Specifically, we show convergence to equilibria, and establish some qualitative facts about the equilibria themselves.

Keywords: Q-learning; prospect theory; cooperative o.d.e.; monotone dynamics; stable equilibria

1. Introduction

Traditional reinforcement learning schemes are concerned with the actions of rational agents seeking to maximize their expected rewards. While these rational agents are risk neutral, reinforcement learning has also been studied under risk-sensitive (risk averse) policies [1, 2, 3]. But according to prospect theory [4] and its sibling *cumulative prospect theory* [5], human beings perceive risk differently in different scenarios: they can be risk seeking in some situations and risk averse in others. (See also [6, 7].) In this work, we study classical Q-learning for Markov decision processes [8], one of the early reinforcement learning algorithms, from a prospect theoretic viewpoint, i.e., when the future returns are distorted using an ‘S-shaped’ valuation map that increases perceived gains and decreases perceived losses.

Previous works [2, 3] applying such prospect theoretic valuation maps worked with certain restrictive assumptions. For example, [2] does not allow for steep valuation maps and high discount factors for future rewards. On the other hand, [3] changes the original Q-learning scheme in a manner that ensures convergence, but the formulation is a departure from the original paradigm as we point out later. In this work, we study the asymptotic behavior of Q-learning scheme when these additional restrictions and/or modifications are dropped. Naturally, we lose global convergence to a single equilibrium, but nevertheless we are able to characterize the asymptotic behavior in qualitative terms to a significant extent. The tools we use are the o.d.e. (for ‘Ordinary Differential Equations’) approach to stochastic approximation [9, 10] (See [11] for a textbook treatment) and monotone dynamical systems [12, 13, 14] to show convergence of the iteration to the set of equilibria and the structure of the latter set.

Our motivation for this study is twofold. The first is the classical motivation behind learning models in economics, viz., to build simple dynamic models to study the qualitative behaviour of boundedly rational macroeconomic agents [15]. Some of the interesting insights we obtain are that the learning in fact equilibrates (i.e., does not get into more complicated behavior such as cycling or worse, strange attractors), though not to a unique equilibrium. Furthermore, the choice of equilibrium depends on the initial condition which is characterizable at least in the cases when they are too high or too low in a certain sense. The equilibration of our model also has a flavor of rational expectations equilibrium [16], which may be an interesting analogy to pursue further.

Secondly, with algorithms taking over from humans in many spheres of human activity, there are ‘human in the loop’ scenarios such as e-commerce, crowdsourced decision making, recommendation networks, etc., where the empirically validated aspects of human judgement such as prospect theory must be factored in. This is so both when not doing so will lead to erroneous predictions and *ipso facto* erroneous decisions, and when the correct outcomes of human peculiarities can lead to undesired outcomes and you want to correct for them. For this, a good theoretical groundwork in terms of mathematical models of behavioral dynamics are important.

^{*}The work of VSB was partially supported by the S. S. Bhatnagar Fellowship from the Government of India.
Email address: borkar.vs@gmail.com, chandak1299@gmail.com (Vivek S. Borkar, Siddharth Chandak)

The need for this is already felt in the rapidly increasing literature that tries to factor in extending prospect theoretic aspects, e.g., in finance [17, 18], game theory [19, 20, 21, 22, 23], the newsvendor problem [24, 25, 26, 27], energy exchange mechanisms [28], etc., the works in prospect theoretic reinforcement learning are relatively few. In addition to those mentioned above, one interesting effort, albeit not in Markov decision theoretic framework, is [29]. See also [30, 31, 32] for other interesting takes on this theme. See [33] and [34] for textbook treatments of prospect theory and reinforcement learning resp., where further pointers to literature in the respective fields can be found.

In the next section, we give a short introduction to Q-learning and briefly, the aspects of prospect theory relevant here. In Section 3, we first present the modified Q-learning scheme and then study its limiting o.d.e. We then use results from monotone dynamical systems to show its convergence to equilibrium points. Then, in Section 4, we show results regarding stability, location and the number of equilibrium points for this scheme. In Section 5, we give a summary of the numerical simulations and comment on our observations. In Section 6, we discuss an alternative modification of the prospect theoretic Q-learning iteration where the total reward is prospect theoretically distorted, not only the future returns.

Notation: For ease of reference, we list the key notation used in the paper in the following table.

Key Notation	
$\{X_n\}$	Controlled Markov chain
$\{U_n\}$	Control process
S	Finite state space with cardinality s
A	Finite action space with cardinality r
α	Discount factor
$k(i, v)$	Reward for choosing action v at state i
k_{min}, k_{max}	min, resp. max of $k(\cdot, \cdot)$
K	The constant $\frac{k_{max}}{1-\alpha}$
$V(i)$	Value function (see (1))
$Q(i, v)$	Q-values (see (2))
$F(\cdot)$	Dynamic programming operator for $Q(\cdot, \cdot)$
$a(n)$	Stepsize sequence
$\mu(i, v, n)$	Number of times action v is chosen at state i till iteration n
$u(\cdot)$	S-shaped map (see Figure 1)
$\xi(i, v)$	i.i.d. noise concentrated in $[-c, c]$
ϵ	Parameter for epsilon-greedy policy (see (13))
$\Lambda(t)$	Diagonal matrix for time scaling under asynchrony

2. Background

2.1. Q-Learning

We sketch the derivation of Watkins' Q-learning algorithm [8] for discounted reward Markov decision processes, as a backdrop for our work. Consider a controlled Markov chain $\{X_n\}$ on a finite state space S , $|S| = s$, governed by a control process $\{U_n\}$ in a finite action space A , $|A| = r$, with controlled transition kernel $(i, j, v) \in S^2 \times A \mapsto p(j|i, v) \in [0, 1]$ with $\sum_j p(j|i, v) = 1$. The controlled Markov property is

$$P(X_{n+1} = i | X_m, U_m, m \leq n) = p(i | X_n, U_n) \quad \forall n.$$

Let $\alpha \in (0, 1)$ be a discount factor and $(i, v) \in S \times A \mapsto k(i, v) \in [0, \infty)$ the per stage reward. For future reference, we denote by k_{min}, k_{max} the minimum and maximum values of $k(i, v)$, which we assume are distinct (i.e., k is not a constant). The infinite horizon discounted reward problem is to maximize

$$E \left[\sum_n \alpha^n k(X_n, U_n) \right]$$

over all $\{U_n\}$ as above, called 'admissible controls'. This maximum for $X_0 = i$ is denoted by $V(i)$. The 'value function' $V(\cdot)$ then satisfies the dynamic programming equation

$$V(i) = \max_v [k(i, v) + \alpha \sum_j p(j|i, v) V(j)], \quad i \in S. \quad (1)$$

V can be computed by, e.g., the value iteration algorithm

$$V_{n+1}(i) = \max_v [k(i, v) + \alpha \sum_j p(j|i, v) V_n(j)], \quad i \in S,$$

beginning with any guess $V_0 \in \mathcal{R}^S$. In reinforcement learning, one seeks a data-driven analog of this, where the nonlinearity due to the ‘max’ operator on the right causes problem because the conditional expectation is inside the nonlinearity.¹ This obstructs any kind of empirical conditional averaging. One way around is to define the Q-values as the expression in square brackets on the right in (1), i.e.,

$$Q(i, v) := k(i, v) + \alpha \sum_j p(j|i, v) V(j), \quad i \in S, v \in A.$$

These satisfy their own dynamic programming equation

$$Q(i, v) = k(i, v) + \alpha \sum_j p(j|i, v) \max_w Q(j, w) \quad \forall i, v. \quad (2)$$

In turn, this can be solved by the ‘Q-value iteration’

$$Q_{n+1}(i, v) = k(i, v) + \alpha \sum_j p(j|i, v) \max_w Q_n(j, w) \quad \forall i, v,$$

where now the conditional expectation w.r.t. $p(\cdot|i, v)$ is outside the max. This facilitates a data driven learning (or stochastic approximation) version as follows. First, when the current state is X_n and the control chosen is U_n , one replaces the conditional expectation on the right hand side by an evaluation at the (real or simulated) next state X_{n+1} , i.e.,

$$Q_{n+1}(i, v) = k(i, v) + \alpha \max_w Q_n(X_{n+1}, w), \quad i = X_n, v = U_n,$$

leaving $Q(j, w), j \neq i$ or $w \neq u$ unchanged. The scheme is then stabilized by making it incremental, i.e., by replacing the full move suggested by the right hand side by a convex combination of it with the previous iterate $Q_n(i, u)$, with a small weight $a(n) \in (0, 1)$ on the former. This leads to the classical Q-learning algorithm

$$\begin{aligned} Q_{n+1}(i, v) &= (1 - a(n)I\{X_n = i, U_n = v\})Q_n(i, v) + \\ &\quad a(n)I\{X_n = i, U_n = v\}(k(i, v) + \alpha \max_w Q_n(X_{n+1}, w)) \\ &= Q_n(i, v) + a(n)I\{X_n = i, U_n = v\} \times \\ &\quad (k(i, v) + \alpha \max_w Q_n(X_{n+1}, w) - Q_n(i, v)). \end{aligned} \quad (3)$$

Here $I\{\dots\}$ is the ‘indicator function’ which is 1 if ‘ \dots ’ holds and 0 otherwise. With the usual Robbins-Monro conditions on the stepsizes $\{a(n)\}$,

$$\sum_n a(n) = \infty, \quad \sum_n a(n)^2 < \infty, \quad (4)$$

this becomes a stochastic approximation algorithm. It is *asynchronous* because only the component corresponding to the current state-action pair is updated at each time. A variant is

$$\begin{aligned} Q_{n+1}(i, v) &= Q_n(i, v) + a(\mu(i, v, n))I\{X_n = i, U_n = v\} \\ &\quad \times (k(i, v) + \alpha \max_w Q_n(X_{n+1}, w) - Q_n(i, v)), \end{aligned} \quad (5)$$

where $\mu(i, v, n) = \sum_{m=0}^n I\{X_m = i, U_m = v\}$ is the ‘local clock’ at (i, v) . Suppose that for some $\delta > 0$,

$$\liminf_{n \uparrow \infty} \frac{\mu(i, v, n)}{n} \geq \delta \quad \forall i, v, \text{ a.s.} \quad (6)$$

¹except in special circumstances that allow a ‘post-decision state’ formulation, see [35].

This ensures ‘sufficient exploration’ in a precise sense, i.e., all state-action pairs are sampled ‘comparably often’. Under some additional restrictions on $\{a(n)\}$ (see Chapter 7, [11]), (5) tracks the ordinary differential equation (ODE)

$$\dot{Q}(t) = F(Q(t)) - Q(t) \quad (7)$$

where $F(x) := [F_{i,v}(x), i \in S, v \in A]$, suitably vectorized, is the dynamic programming operator for Q-values given by

$$F_{i,v}(x) = k(i, v) + \alpha \sum_j p(j|i, v) \max_w x_{j,w} \quad (8)$$

for $x = [x_{i,v}, i \in S, v \in A]$, also suitably vectorized. It can be seen that F is a max-norm contraction:

$$\|F(x) - F(y)\|_\infty \leq \alpha \|x - y\|_\infty,$$

and therefore has a unique fixed point Q^* that satisfies (2). This can be shown to be the globally asymptotically stable equilibrium of (7), from which the a.s. convergence of $\{Q_n\}$ to Q^* can be inferred (see [11], pp. 129-130). A similar analysis also applies to (3) except that the limiting ODE becomes

$$\begin{aligned} \dot{Q}(t) &= \Lambda(t)(F(Q(t)) - Q(t)) \\ &= \tilde{F}_t(Q(t)) - Q(t), \end{aligned}$$

where:

- $\Lambda(t)$ for each t is a diagonal matrix with positive entries on the diagonal that reflect the relative time scaling due to asynchrony ([11], Chapter 7), which can be shown to be bounded away from zero a.s. under (6), and,
- $\tilde{F}_t(x) := (I - \Lambda(t))x + \Lambda(t)F(x)$ which is an $\|\cdot\|_\infty$ -contraction with contraction coefficient $(1 - \delta(1 - \alpha))$ and a unique common fixed point Q^* .

This can be analyzed similarly to (7), but we avoid these complications and stick to (7), because they are not central to our main goals here.

2.2. Prospect Theory

Expected Utility Theory assumes that individuals behave rationally in order to maximize their expected utility. On the other hand, prospect theory [4, 36, 37, 33] aims to describe the actual, empirically validated behavior of people. Prospect theory replaces the utility function with a valuation map over gains and losses defined with respect to a reference point, which is the inflection point of the S-shaped curve in our case. The marginal impact of change in value diminishes with distance from the reference point. Similar to expected utility theory, concavity for gains contributes to risk aversion for gains. On the other hand, convexity for losses contributes to risk seeking behavior.²

We incorporate these ideas into the Q-learning scheme above by passing the estimated future returns through an S-shaped continuous and continuously differentiable map $u(\cdot)$ (see Figure 1). What we take to be ‘estimated future returns’ is, however, a non-unique choice. Some possibilities are:

1. $\max_v Q(X_{n+1}, v)$. *Prima facie*, this seems natural because this is precisely the term that depends on the next state, i.e., the ‘future’. This leads to the scheme

$$\begin{aligned} Q_{n+1}(i, v) &= Q_n(i, v) + a(\mu(i, v, n))I\{X_n = i, U_n = v\} \\ &\quad \times (k(i, v) + \alpha u(\max_w Q_n(X_{n+1}, w)) - Q_n(i, v)). \end{aligned} \quad (9)$$

2. $k(i, v) + \alpha \max_w Q(X_{n+1}, w)$. This is the net estimate of future returns including the immediate reward. This leads to

$$\begin{aligned} Q_{n+1}(i, v) &= Q_n(i, v) + a(\mu(i, v, n))I\{X_n = i, U_n = v\} \times \\ &\quad (u(k(i, v) + \alpha \max_w Q_n(X_{n+1}, w)) - Q_n(i, v)). \end{aligned} \quad (10)$$

²In economics, concave utility function $u(\cdot)$ implies risk aversion because for stochastic returns denoted by X , Jensen’s inequality leads to $E[u(X)] \leq u(E[X])$, meaning ‘sure returns’ $E[X]$ fetch greater utility than random or risky returns X . Analogous statement applies to convex $u(\cdot)$ and risk seeking behavior.




Figure 1: An example of the S-shaped map $u(\cdot)$

3. $k(i, v) + \alpha \max_w Q(X_{n+1}, w) - Q_n(i, v)$. This is the estimate of the incremental future reward. This leads to

$$Q_{n+1}(i, v) = Q_n(i, v) + a(\mu(i, v, n))I\{X_n = i, U_n = v\} \times u(k(i, v) + \alpha \max_w Q_n(X_{n+1}, w) - Q_n(i, v)), \quad (11)$$

This, however, goes against the spirit of the derivation of (5) in the preceding section because the update is no longer a convex combination of the previous iterate and a correction term. Nevertheless, it has been used in the literature [3]. We do not pursue this variant here.

Our interest is in the qualitative analysis of the asymptotic behavior of such algorithms as reflected in the limiting o.d.e. Any of the above leads to a monotone dynamics and our analysis applies, though the actual locations of the equilibria may shift. This is confirmed by our numerical experiments. We mostly focus on the first model.

In fact, we tweak even this model a little by adding noise. This is detailed in the next section.

3. Prospect theoretic Q-learning

3.1. Modified Q-learning scheme

Let $\{X_n\}, \{U_n\}$ be as above. We shall make the additional assumption that the graph of the Markov chain remains irreducible under all control choices at the nodes. Consider the prospect theoretic scheme:

$$Q_{n+1}(i, v) = Q_n(i, v) + a(\mu(i, v, n))I\{X_n = i, U_n = v\} \left(k(i, v) + \alpha u(Q_n(X_{n+1}, U_{n+1}) - \xi_n(X_{n+1}, U_{n+1})) - Q_n(i, v) \right) \quad (12)$$

where $\{\xi_n = [\xi_n(i, v)]\}$ is \mathcal{R}^{sr} -valued zero mean i.i.d. noise. Each $\xi_n(i, v)$ is distributed according to a continuously differentiable density $\varphi(\cdot)$ concentrated on a finite interval $[-c, c]$ for some $c \in [0, k_{min}]$. U_{n+1} is chosen according to an epsilon-greedy policy, i.e., for a prescribed $\epsilon \in (0, 1)$,

$$U_{n+1} = \begin{cases} w_{n+1}^* & \text{w.p. } (1 - \epsilon) \\ w \neq w_{n+1}^* & \text{w.p. } \frac{\epsilon}{r-1} \text{ each} \end{cases} \quad (13)$$

where $w_{n+1}^* = \arg \max_w (Q_n(X_{n+1}, w) - \xi_n(X_{n+1}, w))$, with ties broken by choosing a maximizer with equal probability. We ignore the latter possibility henceforth for sake of simplicity.

The ‘noise’ $\{\xi_n\}$ can be justified as being caused by limited information, noisy measurements, etc. It serves a useful mathematical purpose here, as we note later.

Define $K := \frac{k_{max}}{1-\alpha}$ and let $u : [0, K + c] \mapsto [0, K]$ be continuously differentiable.

Lemma 3.1.1. *If $Q_0 \in \mathcal{S} := [k_{min}, K]^{sr}$, then $Q_n \in \mathcal{S} \forall n \geq 0$.*

Proof. Note that the Q-learning iteration can be written as:

$$\begin{aligned} Q_{n+1}(i, v) &= \left(1 - a(\mu(i, v, n))I\{X_n = i, U_n = v\}\right)Q_n(i, v) \\ &\quad + a(\mu(i, v, n))I\{X_n = i, U_n = v\}\left(k(i, v) + \right. \\ &\quad \left. \alpha u(Q_n(X_{n+1}, U_{n+1}) - \xi_n(X_{n+1}, U_{n+1}))\right). \end{aligned} \quad (14)$$

So $Q_{n+1}(i, v)$ is a convex combination of

$$k(i, v) + \alpha u(Q_n(X_{n+1}, U_{n+1}) - \xi_n(X_{n+1}, U_{n+1}))$$

and $Q_n(i, v)$. Now, if $Q_n \in [k_{min}, K]^{sr}$, then

$$\begin{aligned} k_{min} &\leq k(i, v) + \alpha u(Q_n(X_{n+1}, U_{n+1}) - \xi_n(X_{n+1}, U_{n+1})) \\ &\leq k_{max} + \alpha u(K + c) \\ &= k_{max} + \alpha K \\ &= K. \end{aligned}$$

Hence $Q_n \in \mathcal{S} \Rightarrow Q_{n+1} \in \mathcal{S}$. The claim follows. \square

We assume that the Q-learning iteration was initiated in the set \mathcal{S} . As shown in Chapter 2 of [11], since $u(\cdot)$ is Lipschitz continuous (it is continuously differentiable with bounded first derivative), and $\sup_n \|Q_n\|_\infty \leq K < \infty$, the iteration (12) almost surely tracks the asymptotic behavior of the o.d.e. :

$$\begin{aligned} \frac{d}{dt}q_t(i, v) &= h_{i,v}(q_t) \\ &:= F_{i,v}(q_t) - q_t(i, v) \\ &:= k(i, v) + \\ &\quad \alpha \int_{\mathcal{R}^{sr}} \left(\sum_j p(j|i, v) \left((1 - \epsilon) \max_{w'} (u(q_t(j, w') - y_{j,w})) \right. \right. \\ &\quad \left. \left. + \frac{\epsilon}{r-1} \sum_{w \neq w_{qt,y,j}^*} (u(q_t(j, w) - y_{j,w})) \right) \right) \times \\ &\quad \prod_{j,w} \varphi(y_{j,w}) dy_{j,w} - q_t(i, v). \end{aligned} \quad (15)$$

where $w_{qt,y,j}^* = \arg \max_w (q_t(j, w) - y_{j,w})$. The above implicitly defines the maps $h, F : \mathcal{R}^d \mapsto \mathcal{R}^d$, in fact, $h(x) = F(x) - x$. The integral is a convolution w.r.t. a continuously differentiable function, which makes it continuously differentiable.

Lemma 3.1.2. *When initiated in the set $\mathcal{S} := [k_{min}, K]^{sr}$, o.d.e. (15) stays in the set \mathcal{S} .*

Proof. We can apply the following inequality on the derivative of $q_t(i, v)$:

$$k_{min} - q_t(i, v) \leq \frac{d}{dt}q_t(i, v) \leq k_{max} + \alpha u(K + c) - q_t(i, v)$$

Its discretization can be written as:

$$a_n k_{min} + (1 - a_n)q_n(i, v) \leq q_{n+1}(i, v) \leq a_n K + (1 - a_n)q_n(i, v)$$

This implies that if $q_n(i, v) \in [k_{min}, K]$, then $q_{n+1}(i, v) \in [k_{min}, K]$. So, if initiated in the set $\mathcal{S} := [k_{min}, K]^{sr}$, q_n (and its limit, the o.d.e.) stays in the set \mathcal{S} . \square

The Jacobian matrix of h (resp., F) at q is $J(q) - I$ (resp., $J(q)$), where I is the $sr \times sr$ identity matrix and $J(q)$ for $q \in \mathcal{R}^{sr}$ is the $sr \times sr$ matrix whose $((i, v), (j, w))^{\text{th}}$ element is

$$\begin{aligned} J(q)_{(i,v),(j,w)} &= p(j|i, v) \alpha \int \left[\left((1 - \epsilon) u'(q(j, w) - y_{j,w}) \times \right. \right. \\ &\quad \left. \left. I_{\{q(j,w) - y_{j,w} > q(j,w') - y_{j,w'} \ \forall \ w' \neq w\}} + \frac{\epsilon}{r-1} u'(q(j, w) - y_{j,w}) \right. \right. \\ &\quad \left. \left. \times (1 - I_{\{q(j,w) - y_{j,w} > q(j,w') - y_{j,w'} \ \forall \ w' \neq w\}}) \right) \prod_w \varphi(y_{j,w}) dy_{j,w} \right]. \end{aligned} \quad (16)$$

The continuous differentiability of h was facilitated by the presence of the noise $\{\xi_n\}$, which allows us to exploit the theory of monotone dynamical systems that is available when it holds. For h that are only Lipschitz, such theory appears to be lacking.

3.2. Monotonicity and its consequences

We use the following notion of cooperative o.d.e. from [14]:

Definition 1. (Cooperative o.d.e.) An o.d.e. of the form $\dot{x} = h(x(t))$ is a cooperative o.d.e. if

$$\frac{\partial h_i}{\partial x_j} \geq 0, \quad j \neq i.$$

and the Jacobian matrix for h is irreducible.

Lemma 3.2.1. *When the controlled Markov chain is irreducible, $J(q)$ (the Jacobian of F) is a non-negative irreducible matrix and o.d.e. (15) is a cooperative o.d.e.*

Proof. Since $u' > 0$, it follows that $J(q)$ is a non-negative matrix. $J(q)$ can be written as the product of the following two matrices: a non-negative $sr \times sr$ matrix P where $P_{(i,v),(j,w)} = p(j|i, v)$ and a positive diagonal matrix $J_1(q)$, where $J_1(q)_{(j,w),(j,w)}$ is α times the integral in equation (16)). Since the Markov chain is irreducible, the matrix P is irreducible and hence, the matrix $J(q)$ will be irreducible. Since the off-diagonal terms of $J(q) - I$ are non-negative and the Jacobian is irreducible, (15) is a cooperative o.d.e. \square

Corollary 3.2.0.1. *The dynamical system described by (15) is monotone in the sense that if $q(\cdot), q'(\cdot)$ are two solutions with $q(0) \geq q'(0)$ componentwise, then $q(t) \geq q'(t)$ componentwise $\forall t \geq 0$.*

This follows from the results of [38]. The next theorem follows as a consequence of Theorem 2.1 from [12]

Theorem 3.2.1. *For initial conditions in an open dense set, the solutions of (15) converge to an equilibrium.*

The same must then be true for the iterates of the discrete map $\Phi : \mathcal{S} \mapsto \mathcal{S}$ which maps q_0 to q_1 . Since (15) is cooperative, this map is monotone. It is also ‘order compact’ in the sense of [13], section 5.1, because it maps an order interval, i.e., a set of the form $[y, z]^* := \{x : y \leq x \leq z \text{ componentwise}\}$ to a bounded set. We define $[a, b]^*$ etc. analogously. Using part (b) of Theorem 5.6 from [13], the following holds:

Theorem 3.2.2. *There exist maximal and minimal equilibria q^*, q_* resp. of (15), such that any other equilibrium \hat{q} of (15) satisfies $q_* \leq \hat{q} \leq q^*$ componentwise.*

Since the dynamics preserves order, it follows that for $q_t, t \geq 0$, satisfying (15), $q_0 \geq q^* \implies q_t \geq q^*$ and likewise, $q_0 \leq q_* \implies q_t \leq q_*$. If $q^* > q_*$, since both are fixed points (i.e., equilibria) for the dynamics, $q_* \leq q_0 \leq q^* \implies q_* \leq q_t \leq q^* \forall t \geq 0$ by monotonicity. Since the map $q_0 \mapsto q_1$ is continuous, the following theorem and its corollary hold:

Theorem 3.2.3. (Order Interval Trichotomy, Theorem 5.1 from [13]) *At least one of the following holds:*

1. \exists a third equilibrium \hat{q} such that $q_* < \hat{q} < q^*$,
2. \exists a trajectory q_t of (15) such that $q_t \uparrow q^*$ as $t \uparrow \infty$ and $q_t \downarrow q_*$ as $t \downarrow -\infty$,
3. \exists a trajectory q_t of (15) such that $q_t \downarrow q_*$ as $t \uparrow \infty$ and $q_t \uparrow q^*$ as $t \downarrow -\infty$.

Corollary 3.2.3.1. (Corollary 5.2 from [13]) *If both q_* and q^* are stable, then there is at least one more equilibrium \hat{q} such that $q_* < \hat{q} < q^*$.*

4. Equilibrium Points

4.1. Regions with stable equilibria

As our discrete map is the time-1 map of a differential equation with smooth right hand side, the stability of its equilibria, which are the same as equilibria of the differential equation, can be analyzed by looking at the linearization of $h(\cdot) := [h_{i,v}(\cdot)]_{i \in S, v \in A}$ at the equilibrium, i.e., the eigenvalues of the Jacobian matrix $J(q) - I$ evaluated at the equilibrium. We have the following interesting observation from the foregoing.

Theorem 4.1.1. *If q^* , resp. q_* , is hyperbolic, then it is stable.*

Proof. If (say) q^* is hyperbolic, it is isolated by the inverse function theorem and has well defined stable and unstable manifolds in an open neighborhood, corresponding to the eigenvalues of $J(q^*) - I$ in the left, resp. right half of the complex plane. Furthermore, if the dimension of the unstable manifold is ≥ 1 (i.e., there is at least one unstable eigenvalue), then for all initial conditions in an open neighborhood O of q^* , the trajectories that are not initiated exactly on the stable manifold eventually move away from q^* . Since the dimension of the stable manifold is strictly less than s , this is so for initial conditions in an open dense subset O' of O . We may further exclude from O' its intersection with the closed nowhere dense set where the convergence claim of Theorem 3.2.1 can fail and denote the resultant open dense subset of O as O'' . Denote the intersection of O'' with the open cone $C^* := \{q : q = q^* + y, y > 0 \text{ componentwise}\}$, as \hat{O} , which will be an open set. Then $q(0) \in \hat{O}$ implies that $q(t)$ eventually moves away from q^* . By monotonicity, $q(0) \geq q^*$ componentwise $\implies q(t) \geq q^* \forall t \geq 0$ componentwise by monotonicity. Since q^* is the maximal equilibrium, if $q(0) \in \hat{O}$, it must hold that $q(t) \rightarrow q^*$, a contradiction to the previous claim. Hence q^* must be stable. The claim for q_* is proved similarly. \square

The following theorem gives bounds on the eigenvalues of $J(q) - I$:

Theorem 4.1.2. (Perron-Frobenius Theorem [39]) *Let A be a square non-negative irreducible matrix. Then*

1. *A has a real positive eigenvalue λ_A which is greater than or equal to the absolute value of any other eigenvalue of A .*
2. *$r \leq \lambda_A \leq R$ where $r = \min_i r_i$ and $R = \max_i r_i$, where r_i denotes the sum of the elements of row i of A .*

Let the sum of the $(i, v)^{\text{th}}$ row of $J(q)$ be $\Gamma(q)_{i,v}$. Let $\Gamma(q)^* = \max_{i,v} \Gamma(q)_{i,v}$ and $\Gamma(q)_* = \min_{i,v} \Gamma(q)_{i,v}$. Since $J(q)$ is a non-negative irreducible matrix, there exists a real positive eigenvalue λ^* (dominant eigenvalue) for $J(q)$ such that any other eigenvalue λ of $J(q)$ has its absolute value (and hence its real part) smaller than or equal to λ^* , by Theorem 4.1.2. We also know that $\Gamma(q)_* \leq \lambda^* \leq \Gamma(q)^*$. For any eigenvalue λ of $J(q)$, $\lambda - 1$ is an eigenvalue of the Jacobian $J(q) - I$. So, the real part of all eigenvalues of $J(q) - 1$ are less than $\lambda^* - 1$.

Since $u(\cdot)$ is an S-shaped function, we know that $u'(x)$ is typically small for sufficiently low or high values of x . Thus $u'(x) < 1 < \frac{1}{\alpha}$ for sufficiently low or high values of x and can exceed $\frac{1}{\alpha}$ in the mid-range. If $u'(x) < \frac{1}{\alpha} \forall x \in [0, K + c]$, then we can use the results from [2] which show that there will exist only one equilibrium point in the set and it will be stable. In fact this follows easily from the contraction mapping theorem. We consider here the case where $u'(x)$ exceeds $\frac{1}{\alpha}$ in the middle region. Define points $a, b, b > a$, in $[0, K]$ as the largest and smallest points, respectively, in $[0, K]$ such that $u'(x) < \frac{1}{\alpha} \forall x \in [0, a) \cup (b, K + c]$.

Theorem 4.1.3. *There is at most one equilibrium point for (15) in the set $(b + c, K]^{sr}$. If such an equilibrium point exists, it will be a stable equilibrium and the maximal equilibrium point. Similarly, there is at most one equilibrium point for (15) in the set $[k_{\min}, a - c)^{sr}$ and if such an equilibrium point exists, it will be a stable equilibrium and the minimal equilibrium point.*

Proof. Let q_1 be any point in the set $(b + c, K]^{sr}$. Then $[q_1 - c, q_1 + c]^* \subseteq (b, K + c]$ componentwise. Since $u'(\cdot) < 1/\alpha$ in this region, $\Gamma(q_1)_{i,v} < 1, \forall i, v$ and hence, $\lambda^* < 1$. Hence real parts of all eigenvalues of the Jacobian $J(q_1) - I$ are negative. So any equilibrium point lying in this region will be hyperbolic and stable. Now suppose that there are two equilibria q_1, q_2 in the aforementioned set. The two points can be ordered or unordered. First we consider the case where they are ordered and $q_1 < q_2$. By Corollary 3.2.3.1, there exists another equilibrium point q_3 , such that $q_1 < q_3 < q_2$. Then q_3 will also be a stable equilibrium and hence there will be more stable equilibrium points between q_1, q_3 , and between q_3, q_2 .

Repeated application of this argument implies that we will have a continuum of non-isolated equilibria. But real part of all eigenvalues of the Jacobian $J(q) - I$ are negative in this region, implying that all equilibria are isolated. This gives us a contradiction. Hence there cannot be two ordered equilibria in the region.

Now consider the other case in which there are two unordered equilibria q_1, q_2 in the region $(b+c, K]^{sr}$. We know that there exists q^* such that all equilibrium points q satisfy $q \leq q^*$ (Theorem 3.2.2). Since no ordering exists between q_1 and q_2 , they can't be equal to q^* . So, $q_1 < q^*$ where both q_1 and q^* lie in this region. But we have shown earlier that there cannot exist ordered equilibria in the region. So, at most one equilibrium point can exist in this region and that will be the maximal equilibrium point. Analogous statement for the set $[k_{min}, a-c]^{sr}$ is proved similarly. \square

Based on Theorem 4.1.3, we subsequently refer to the sets $[k_{min}, a-c]^{sr}$ and $(b+c, K]^{sr}$ as the lower and upper stable regions, respectively.

4.2. Additional results on stability of equilibria

Let points d, e in $[0, K]$ be the smallest and the largest points in $[0, K]$ such that $u'(x) > \frac{1}{\alpha} \forall x, d < x < e$. In most cases, d and e will be same as a and b respectively, but we have defined them separately to take care of cases where $u'(x) = \frac{1}{\alpha}$ in the intervals $x \in [a, d]$ and $x \in [e, b]$.

Theorem 4.2.1. *Any equilibrium point in the region $(d+c, e-c)^{sr}$ is an unstable equilibrium point.*

Proof. Let q_1 be any point in the set $(d+c, e-c)^{sr}$. Then, $[q_1-c, q_1+c]^* \subseteq (d, e)$ componentwise. Since $u'(\cdot) > 1/\alpha$ in this region, $\Gamma(q_1)_{i,v} > 1, \forall i, v$ and hence, $\lambda^* > 1$. Thus at least one eigenvalue of $J(q) - I$ has a positive real part. Therefore any equilibrium point in this region will be unstable. \square

Let the maximum value of $u'(x)$ be attained at m_1 .

Theorem 4.2.2. *If all equilibrium points are hyperbolic and $u(x)$ is convex, respectively, concave in the regions $x < m_1$ and $x > m_1$, respectively, then there can exist at most one stable equilibrium point in the region $[k_{min}, m_1-c]^{sr}$. Similarly in the region $(m_1+c, K]^{sr}$, there can exist at most one stable equilibrium. If these exist, then they will be the minimal and maximal equilibrium points, respectively.*

Proof. Suppose there exist two stable equilibrium points q_1, q_2 in the region $(m_1+c, K]^{sr}$. Again, there are two possibilities: they can be ordered or unordered. Let us first consider the case where they are ordered, say $q_1 < q_2$. Since we have taken $u(\cdot)$ to be concave in the region, for any point $q_3 \in [q_1, q_2]^*$, $u'(q_1(i, v) - y) \geq u'(q_3(i, v) - y) \geq u'(q_2(i, v) - y) \forall y \in [-c, c]$. Hence $\Gamma(q_1) \geq \Gamma(q_3) \geq \Gamma(q_2)$ componentwise. Let $\lambda_{q_1}^*, \lambda_{q_2}^*$ and $\lambda_{q_3}^*$ be the dominant eigenvalues of $J(q_1), J(q_2)$ and $J(q_3)$ respectively. Then $\Gamma(q_1) \geq \Gamma(q_3) \geq \Gamma(q_2)$ implies that $\lambda_{q_1}^* \geq \lambda_{q_3}^* \geq \lambda_{q_2}^*$ (Theorem A.9 from [39]). Since q_1 and q_2 are stable equilibria, $\lambda_{q_1}^*$ and $\lambda_{q_2}^*$ are less than 1, and hence $\lambda_{q_3}^* < 1$. So, q_3 will also be a stable equilibrium. As shown in the proof of Theorem 4.1.3, this gives us a continuum of non-isolated equilibria, which contradicts our assumption that all equilibria are hyperbolic.

If there exist two unordered equilibrium points in the region, then there will exist another equilibrium point q_3 which will be the maximal equilibrium and will be stable, as we have assumed that all equilibria are hyperbolic. Applying the first part of this proof to q_1 and q_3 gives us a contradiction. Hence there can be at most one stable equilibrium point in the region $(m_1+c, K]^{sr}$. The same is true for the set $[k_{min}, m_1-c]^{sr}$. \square

This theorem can also be applied where the valuation map is a traditional utility function which is either convex or concave in the whole domain. In our case, however, there can exist many other stable equilibrium points with some components below and some above m_1 .

Let b_1 be a point such that $k_{min} + \alpha u(b_1 - c) \geq b_1$. Then $k_{min} + \alpha u(x - c) \geq b_1 \forall x \geq b_1$. In the convex combination form of the iteration (14), note that if $Q_n \geq b_1$, then

$$\begin{aligned} k(i, v) + \alpha u(Q_n(X_{n+1}, U_{n+1}) - \xi_n(X_{n+1}, U_{n+1})) \\ \geq k_{min} + \alpha u(b_1 - c) \geq b_1 \end{aligned}$$

and hence $Q_{n+1} \geq b_1$. Similar to Lemma 3.1.1, we note that if the iteration (12) is initiated in the set $[b_1, K]^{sr}$, then it will stay in this set. In the following theorem, we use this to provide a sufficient condition for a stable equilibrium point of (15) to exist in the upper stable region. For simplicity, define $u_1(x) := k_{min} + \alpha u(x - c)$.

Theorem 4.2.3. *If $u_1(b+c) \geq b+c$, then there exists a stable maximal equilibrium point of (15) in the region $[b+c, K]^{sr}$ and when the iteration (12) is initiated in this set, it converges to this equilibrium point.*

Proof. In the region $(b+c, K]$, $u'_1(x) < 1$ and hence u_1 and $y = x$ can intersect at most once in the region. Such an intersection point would exist if $u_1(b+c) > b+c$. If $u_1(b+c) = b+c$, then that point will be $b+c$ and there will be no other intersection. Let this intersection point be called g .

As stated before, when $u_1(b+c) \geq b+c$, then the iteration (12) would stay in the set $[b+c, K]^{sr}$ when initiated in it. Similarly, if initiated in the set $[g, K]^{sr}$, the iteration stays in that set. Hence there exists an equilibrium point of (15) in this set. As proved in Theorem 4.1.3, there can be at most one equilibrium point in the set $[b+c, K]^{sr}$. Hence, if $u_1(b+c) \geq b+c$ the equilibrium point in the set $[g, K]^{sr}$ will be the stable maximal equilibrium point and whenever iteration (12) is initiated in the region $[b+c, K]^{sr}$, it will converge to this equilibrium point. \square

Similar to $u_1(\cdot)$, define $u_2(x) := k_{max} + \alpha u(x+c)$. We have the following theorem relating $u_2(\cdot)$ and the lower stable region similar to Theorem 4.2.3:

Theorem 4.2.4. *If $u_2(a-c) \leq a-c$, then there exists a stable minimal equilibrium point of (15) in the region $[k_{min}, a-c]^{sr}$ and when the iteration (12) is initiated in this set, it converges to this equilibrium point.*

Theorem 4.2.5. *If $u_1(a+c) > a+c$, then there exists only one equilibrium point of (15) in the set $[k_{min}, K]^{sr}$ and it will lie in the region $(b+c, K]^{sr}$.*

Proof. We know that $u'_1(x) = \alpha u'(x-c) \geq 1, \forall x \in [a+c, b+c]$ (by the definition of a, b) and hence $u_1(a+c) > a+c$ implies that $u_1(b+c) > b+c$. From the proof of Theorem 4.2.3, we know that the curve $u_1(x)$ and $y = x$ intersect only once in the region $(b+c, K]$. Let this point of intersection be g . For all $x \in [a+c, b+c]$, $u'_1(x) \geq 1$ and hence,

$$u_1(x) - u_1(a+c) \geq x - (a+c) \implies u_1(x) \geq u_1(a+c) + x - (a+c) > x.$$

For all $x \in [k_{min}, a+c]$, $u'_1(x) \leq 1$ and hence,

$$u_1(a+c) - u_1(x) \leq (a+c) - x \implies u_1(x) \geq u_1(a+c) - (a+c) + x > x.$$

So $u_1(x)$ lies strictly above $y = x$ in the region $[k_{min}, b+c]$ and hence g will be the only intersection of these two curves in the region $[k_{min}, K]$. Consider any $q \notin [g, K]^{sr}$. Then its minimum component is in $[k_{min}, g)$. Let this minimum component be $q_{i',v'}$ (i.e., $(i', v') = \arg \min_{(i,v)} q_{i,v}$). Since $u_1(q_{i',v'}) > q_{i',v'}$ and $q_{i,v} \geq q_{i',v'}, \forall (i, v)$,

$$\begin{aligned} h_{i',v'}(q) &\geq k_{i',v'} + \alpha u(q_{i',v'} - c) - q_{i',v} \\ &\geq k_{min} + \alpha u(q_{i',v'} - c) - q_{i',v'} = u_1(q_{i',v'}) - q_{i',v'} \\ &> 0. \end{aligned}$$

So there can be no equilibrium point of (15) which does not belong to $[g, K]^{sr}$. By Theorem 4.2.3, we know that there exists one stable equilibrium in the region $[g, K]^{sr}$. This will be the only equilibrium point and will be both the maximal and minimal equilibrium point of (15). \square

5. Numerical Experiments

We numerically simulated our modified Q-learning scheme and the corresponding o.d.e. to verify our results and gain additional insights. We used the shifted and scaled version of sigmoid logistic function:

$$u(x) = \frac{L}{1 + e^{-\gamma(x-x_0)}}.$$

By modifying value of L, γ and x_0 , we explored the behavior of the Q-learning scheme with varying maximum value, steepness and midpoint of the curve, respectively. Both very small and large state and action spaces were explored with values of both s and r ranging from 2 to 100. To satisfy (4), we chose $a(\mu(i, v, n))$ as:

$$a(\mu(i, v, n)) = \frac{1}{\left\lceil \frac{\mu(i, v, n)}{100} \right\rceil}$$

where $\lceil \cdot \rceil$ is the ceiling function and the scaling by 100 was chosen for faster convergence. The rewards k were generated randomly in a given range set by fixing k_{min} and k_{max} . The transition matrix was generated randomly. We tried different values of the discount factor α ranging from 0.01 to 0.99. We used the raised cosine distribution for noise as we required the noise distribution to be continuously differentiable in a finite support. Noise was in general kept much smaller than the rewards ($c \approx 0.01$). Finally, we used $\epsilon = 0.05$ for the ϵ -greedy scheme.




Figure 2: Convergence plot of modified Q-learning scheme: (a) shows the Bellman error for modified Q-learning scheme, while (b) shows the moving average of the same over 1000 iterations. ($\alpha = 0.5, k_{min} = 2, k_{max} = 5, s = r = 20, \epsilon = 0.05$ and iteration initiated in $[K - 1, K - 0.5]^{sr}$)

We varied different parameters and studied the Q-learning iteration, the o.d.e. and their equilibrium points. These are some of our observations:

- As long as the conditions mentioned in the Section 3 are met, the Q-learning iteration and the o.d.e. converged to an equilibrium point and to the same point when initiated at the same point. In Figure 2, we plot a representative convergence plot of error vs. iterations where error for iteration n is defined as: $|Q_{n+1}(X_n, U_n) - Q_n(X_n, U_n)|$.
- Choice of $\{a(n)\}$ and ϵ have an impact on the rate of convergence, but do not observably affect the equilibrium points.
- As expected, when either α is too small or the function $u(\cdot)$ rises very gradually (i.e., $u'(x) < \frac{1}{\alpha}$ in the whole region), then there exists only one equilibrium point (see Figure 3(a)). Even when $u'(x)$ exceeds $\frac{1}{\alpha}$ in the middle region, we frequently observe only one equilibrium point in \mathcal{S} when the maximum value of the first derivative of $u(\cdot)$ is not very high.




Figure 3: $u(\cdot)$ and $u'(\cdot)$ for two different cases: (a) shows the case where $u(\cdot)$ rises gradually and $\max_x u'(x) < 1/\alpha$, ($\alpha = 0.2$). Only one equilibrium is observed in this case. (b) shows the case where $\max_x u'(x) \geq 1/\alpha$, ($\alpha = 0.6$). Equilibria are observed in both lower and upper stable regions in this case. ($k_{min} = 2, k_{max} = 5, s = r = 20, \epsilon = 0.05$)

- For $u(\cdot)$ that is very steep at the point of inflection, the iteration usually converges to one of the

two equilibria, one each in the upper and lower zones, depending on the initiation. (see Figure 3(b))

- Theorems 4.2.3 and 4.2.4 only provide sufficient (and not necessary) conditions for existence of equilibria in the upper and lower stable regions. And hence, even when $u_1(b+c) < b+c$, we frequently observe an equilibrium point in the upper stable region (see Figure 4).




Figure 4: $u_1(\cdot)$ in two different cases: (a) shows the case where $u_1(b+c) \geq b+c$, ($k_{min} = 10$). (b) shows the case where $u_1(b+c) < b+c$, ($k_{min} = 1$). In both cases, an equilibrium exists in the upper stable region. ($\alpha = 0.5, k_{max} = 20, s = r = 20, \epsilon = 0.05$)

5.1. Third stable equilibrium point

In general, we observed that convergence to any equilibrium point other than the maximal and minimal ones in the upper and lower stable regions is very rare. In our initial experiments, we noticed that the iteration converged either to the maximal or to the minimal equilibrium point. To confirm that the existence of a third stable equilibrium point is possible, we manually constructed and computed the equilibrium points for a small system ($s = 4, r = 2$). For further simplicity, we kept all rewards very close to 2 (i.e., $k(i, v) \in [1.99, 2.01], \forall i, v$) and fixed the value of α to 0.8. We chose a very steep $u(\cdot)$ such that $u(x) \approx 0, \forall x \in [k_{min}, a]$ and $u(x) \approx K = 10, \forall x \in (b, K]$. The function $u(\cdot)$ rises from 0 to 10 almost entirely between a and b (see Figure 5). This also ensured that the conditions in Theorems 4.2.3 and 4.2.4 are satisfied. The transition matrix was fixed so that for each state, the two actions are identical (i.e. $p(j|i, v) = p(j|i, w), \forall i, j$ where w, v are the two actions for state i). For the transition probabilities we chose, we observed that there are four stable equilibria (two equilibria in addition to the maximal and minimal equilibria). When initiated in close vicinity to these additional equilibrium points, the iteration converges to them.




Figure 5: $u(\cdot)$ with a very steep rise constructed in order to confirm existence of a third stable equilibrium point

Apart from this above constructed case, we never observed the Q-learning iteration to converge to a stable equilibrium point other than the maximal or minimal. While three or more stable equilibria

can exist for many systems, convergence to these points seems very infrequent in randomly generated test cases, suggesting a small domain of attraction. This is a purely empirical observation and it will be interesting to see if there is a fundamental reason for it being so.

We also compared our scheme with classical Q-learning for the same problem parameters (figures not included). The (necessarily unique) equilibrium of the latter was empirically found to be closer to the maximal equilibrium of the former than its minimal equilibrium, and either above or below it depending on the shape of the S-curve. For example, the equilibrium of classical Q-learning would be above the maximal equilibrium of our scheme if the S-curve had a very limited range. It became closer to the minimal equilibrium if we replaced reward maximization by cost minimization. Though the latter is merely a mathematical curiosity, it suggests that the preceding observations have more to do with ours being a maximization problem.

6. Alternative Formulation

In our original formulation, only the future returns are distorted using the prospect theoretic valuation map. We studied another formulation where the S-shaped curve $u(\cdot)$ is applied to the total returns, i.e., both the current rewards and the future returns are distorted.

6.1. Q-learning scheme and its convergence

The Q-learning iteration in this case is the following:

$$\begin{aligned} Q_{n+1}(i, v) &= Q_n(i, v) + a(\mu(i, v, n)) I\{X_n = i, U_n = v\} \times \\ &\quad \left(u\left(k(i, v) + \alpha(Q_n(X_{n+1}, U_{n+1}) - \right. \right. \\ &\quad \left. \left. \xi_n(X_{n+1}, U_{n+1}))\right) - Q_n(i, v) \right). \end{aligned} \quad (17)$$

In this case, $u : [0, K + \alpha c] \mapsto [0, K]$ and similar to Lemma 3.1.1, when initiated in the set $\mathcal{S}_1 := [0, K]^{sr}$, the Q-learning iteration stays in the set \mathcal{S}_1 .

Iteration (17) tracks to the following o.d.e.:

$$\begin{aligned} \frac{d}{dt} q_t(i, v) &= h_{1_{i,v}}(q_t) \\ &= F_{1_{i,v}}(q_t) - q_t(i, v) \\ &:= \int_{\mathcal{R}^{sr}} \left[\left(\sum_j p(j|i, v) \left((1 - \epsilon) u\left(k(i, v) \right. \right. \right. \right. \\ &\quad \left. \left. \left. + \alpha \max_w (q_t(j, w) - y_{j,w}) \right) \right) \right. \\ &\quad \left. + \frac{\epsilon}{r-1} \sum_{w \neq w_{q_t^*, v, j}^*} u\left(k(i, v) + \alpha(q_t(j, w) - y_{j,w})\right) \right) \right) \\ &\quad \times \prod_{j,w} \varphi(y_{j,w}) dy_{j,w} \Big] - q_t(i, v). \end{aligned} \quad (18)$$

As in Lemma 3.1.2 above, o.d.e. (18) remains in the set \mathcal{S}_1 when initiated in \mathcal{S}_1 . This o.d.e. is also a cooperative o.d.e. where the Jacobian matrix of h_1 (resp., F_1) at q is $J_1(q) - I$ (resp., $J_1(q)$) with the $((i, v), j, w)^{\text{th}}$ element of $J_1(q)$ being:

$$\begin{aligned} J_1(q)_{(i,v),(j,w)} &= p(j|i, v) \alpha \\ &\quad \times \int \left[\left((1 - \epsilon) u'\left(k(i, v) + \alpha(q(j, w) - y_{j,w})\right) I_\beta(j, w) \right. \right. \\ &\quad \left. \left. + \frac{\epsilon}{r-1} u'\left(k(i, v) + \alpha(q(j, w) - y_{j,w})\right) (1 - I_\beta(j, w)) \right) \right. \\ &\quad \left. \times \prod_w \varphi(y_{j,w}) dy_{j,w} \right] \end{aligned} \quad (19)$$

where $I_\beta(j, w) = I_{\{q(j, w) - y_{j, w} > q(j, w') - y_{j, w'} \mid w' \neq w\}}$. (As before, we ignore the case of multiple maximizers.)

Since o.d.e. (18) is a cooperative o.d.e. and it maps \mathcal{S}_1 to \mathcal{S}_1 , all theorems in Section 3.2 can be applied to this alternative formulation as well. Thus there exist maximal and minimal equilibria for (17) and the order interval trichotomy holds for the equilibrium points of (17).

6.2. Stable regions

When $u'(x) < 1/\alpha, \forall x \in \mathcal{S}_1$, then the results from [2] again show that there will exist only one equilibrium point in the set \mathcal{S}_1 . But when $u'(x)$ crosses $1/\alpha$ in the middle regions, we observe behavior different from our original formulation. The upper and lower stable regions do not always exist in this case. As in Section 4.1, these stable regions are defined in the regions where the sum of each row of the Jacobian matrix $J_1(q)$ is less than 1. This will be true when $u'(k(i, v) + \alpha(q(i, v) \pm c)) < \frac{1}{\alpha}, \forall i, v$.

So, the upper stable region is defined as $(b' + c, K]^{sr}$ where $b' = \frac{b - k_{min}}{\alpha}$ and will exist if the following condition holds:

Condition 1.

$$b' + c < K \Leftrightarrow \frac{b - k_{min}}{\alpha} + c < K \Leftrightarrow b < k_{min} + \alpha(K - c).$$

Similarly the lower stable region is defined as $[0, a' - c)^{sr}$ where $a' = \frac{a - k_{max}}{\alpha}$ and will exist if:

Condition 2.

$$a' - c > 0 \Leftrightarrow \frac{a - k_{max}}{\alpha} - c > 0 \Leftrightarrow a > k_{max} + \alpha c.$$

6.3. Numerical experiments

To simulate this formulation, i.e., iteration (17) and o.d.e. (18), we vary parameters as in Section 5. As expected, this scheme also converges to the equilibrium points. Most trends observed for the original scheme are observed here as well. For example, when $u(\cdot)$ rises very steeply and α is large (greater than ≈ 0.7), we observe two equilibrium points, one each in the upper and lower stable regions.

An important difference between the two schemes lies in the values of the maximal and minimal equilibrium points. Since both the current and the future returns are distorted in this scheme, when the original scheme converged to large values of Q , this scheme converges to even larger Q values. Similarly, when the original scheme converged to small Q values, this scheme converges to even smaller Q values. This is observed in the simulations as well. When same parameters are set for both schemes, the maximal equilibrium point of the alternate formulation is higher than the maximal equilibrium for the original formulation.

7. Conclusions

In this work we studied classical Q-learning from a prospect theoretic viewpoint, i.e., when the valuation of future returns is distorted by an S-shaped subjective map. We then present conditions under which the iteration and its limiting o.d.e. converge to equilibrium points. Upper and lower stable regions are defined, in each of which at most one equilibrium exists and is stable. Additional results regarding the number and location of equilibria are also presented. We verify these results through simulations and make further comments regarding the observations. We finally study an alternative prospect theoretic scheme where both the current and future returns are distorted. A possible avenue for future work is to characterize conditions under which there are three or more equilibria and the scheme converges to an equilibrium point other than the maximal or minimal equilibrium. Also, it will be interesting to study the bifurcation phenomena that arise when the map $u(\cdot)$ is homotopically morphed from the identity map to an S-curve.

References

- [1] Mihatsch, O. and Neuneier, R., 2002. Risk-sensitive reinforcement learning. *Machine Learning* 49(2-3), 267-290.
- [2] Shen, Y., Stannat, W. and Obermayer, K., 2013. Risk-sensitive Markov control processes. *SIAM Journal on Control and Optimization* 51(5), 3652-3672.

- [3] Shen, Y., Tobia, M. J., Sommer, T. and Obermayer, K., 2014. Risk-sensitive reinforcement learning. *Neural Computation* 26(7), 1298-1328.
- [4] Kahneman, D. and Tversky, A., 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, 99-127.
- [5] Tversky, A. and Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297-323.
- [6] Dacey, R., 2003. The S-Shaped Utility Function. *Synthese* 135, 243-272.
- [7] Armstrong, J., and Brigo, D., 2019. Risk managing tail-risk seekers: VaR and expected shortfall vs S-shaped utility. *Journal of Banking & Finance*, 101, 122-135.
- [8] Watkins, C. J. C. H., 1989. Learning from delayed rewards. *Ph.D. Thesis*, King's College, University of Cambridge, UK.
- [9] Derevitskii, D. P. and Fradkov, A. L., 1974 Two models for analyzing the dynamics of adaptation algorithms'. *Automation and Remote Control* 35, 59-67.
- [10] Ljung, L., 1977. Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control* 22, 551-575.
- [11] Borkar, V. S., 2008. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Hindustan Publishing Agency, New Delhi, and Cambridge University Press, Cambridge, UK.
- [12] Hirsch, M. W. and Smith, H. L., 2003. Monotone systems, a mini-review, in *Positive Systems* (Benvenuti L., De Santis A. and L. Farina (eds)), Lecture notes on Control and Information sciences, Springer, Berlin-Heidelberg, 183-190.
- [13] Hirsch, M. W. and Smith, H., 2005. Monotone maps: a review. *Journal of Difference Equations and Applications* 11(4-5), 379-398.
- [14] Smith, H. L., 2008. *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, Providence, R. I.
- [15] Fudenberg, D. F. and Levine, D. K., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- [16] Begg, D. K. H., 1982. *The Rational Expectations Revolution in Macroeconomics: Theories & Evidence*. The Johns Hopkins University Press, Baltimore, MD.
- [17] He, X. D. and Zhou, X. Y., 2011. Portfolio choice under cumulative prospect theory: An analytical treatment. *Management Science* 57(2), 315-331.
- [18] Zhang, L., Zhang, H., and Hao, S., 2018. An equity fund recommendation system by combing transfer learning and the utility function of the prospect theory. *The Journal of Finance and Data Science*, 4(4), 223-233.
- [19] Phade, S. R. and Anantharam, V., 2018. Learning in Games with Cumulative Prospect Theoretic Preferences. *arXiv preprint arXiv:1804.08005*.
- [20] Phade, S. R. and Anantharam, V., 2019. On the geometry of Nash and correlated equilibria with cumulative prospect theoretic preferences. *Decision Analysis* 16(2), 142-156.
- [21] Phade, S. R. and Anantharam, V., 2020. Black-Box Strategies and Equilibrium for Games with Cumulative Prospect Theoretic Players. *arXiv preprint arXiv:2004.09592*.
- [22] Hota, A. R., Garg, S. and Sundaram, S., 2016. Fragility of the commons under prospect-theoretic risk attitudes. *Games and Economic Behavior*, 98, 135-164.
- [23] Tian, R., Sun, L. and Tomizuka, M., 2021. Bounded Risk-Sensitive Markov Games: Forward Policy Design and Inverse Reward Learning with Iterative Reasoning and Cumulative Prospect Theory. In *AAAI Conference on Artificial Intelligence*.

- [24] Vipin, B., and Amit, R. K., 2019. Describing decision bias in the newsvendor problem: A prospect theory model. *Omega*, 82, 132-141.
- [25] Nagarajan, M., and Shechter, S., 2014. Prospect Theory and the Newsvendor Problem. *Management Science*, 60(4), 1057-1062.
- [26] Shen, Y., Zhao, X., and Xie, J., 2017. Revisiting prospect theory and the newsvendor problem. *Operations Research Letters*, 45(6), 647-651.
- [27] Surti, C., Celani, A., and Gajpal, Y., 2020. The Newsvendor Problem: The Role of Prospect Theory and Feedback. *European Journal of Operational Research*, 287.
- [28] Xiao, L., Mandayam, N. B. and Poor, H. V., 2014. Prospect Theoretic Analysis of Energy Exchange Among Microgrids. *IEEE Transactions on Smart Grid*, 6(1), 63-72.
- [29] Prashanth, L. A., Jie, C., Fu, M., Marcus, S. and Szepesvári, C., 2016. Cumulative prospect theory meets reinforcement learning: Prediction and control. In *Proc. International Conference on Machine Learning*, 1406-1415.
- [30] Denrell, J. C., 2008. Reinforcement Learning Leads to Risk Averse Behavior. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 30, No. 30).
- [31] Feldmaier, J., Meyer, D., Shen, H., and Diepold, K., 2015. Reinforcement learning with preferences. In *The 2nd Multidisciplinary Conference on Reinforcement Learning and Decision Making* (pp. 143-147).
- [32] Ratliff, L. J., and Mazumdar, E., 2019. Inverse Risk-Sensitive Reinforcement Learning. *IEEE Transactions on Automatic Control*, 65(3), 1256-1263.
- [33] Wakker, P. P., 2010. *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, Cambridge, UK.
- [34] Bertsekas, d. P., 2019. *Reinforcement Learning and Optimal Control*. Athena Scientific, Belmont, MA.
- [35] Powell, W. B., 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality* (2nd ed.). John Wiley and Sons, New York.
- [36] Fox, C. R. and Poldrack, R. A., 2009. Prospect theory and the brain. In *Neuroeconomics*, Academic Press, 145-173.
- [37] Ruggeri, K., Alí, S., Berge, M.L., Bertoldo, G., Bjørndal, L.D., Cortijos-Bernabeu, A., Davison, C., Damić, E., Esteban-Serna, C., Friedemann, M. and Gibson, S.P., 2020. Replicating patterns of prospect theory for decision under risk. *Nature Human Behaviour*, 1-12.
- [38] Hirsch, M. W., 1985. Systems of differential equations that are competitive or cooperative II: Convergence almost everywhere. *SIAM Journal on Mathematical Analysis* 16(3), 423-439.
- [39] Abramov, A. P., *Balanced and Cyclical Growth in Models of Decentralized Economy*, Springer International Publishing Switzerland 2014