
Decentralized Reinforcement Learning: Global Decision-Making via Local Economic Transactions

Michael Chang¹ Sidhant Kaushik¹ S. Matthew Weinberg² Thomas L. Griffiths² Sergey Levine¹

Abstract

This paper¹ seeks to establish a framework for directing a society of simple, specialized, self-interested agents to solve what traditionally are posed as monolithic single-agent sequential decision problems. What makes it challenging to use a decentralized approach to collectively optimize a central objective is the difficulty in characterizing the equilibrium strategy profile of non-cooperative games. To overcome this challenge, we design a mechanism for defining the learning environment of each agent for which we know that the optimal solution for the global objective coincides with a Nash equilibrium strategy profile of the agents optimizing their own local objectives. The society functions as an economy of agents that learn the credit assignment process itself by buying and selling to each other the right to operate on the environment state. We derive a class of decentralized reinforcement learning algorithms that are broadly applicable not only to standard reinforcement learning but also for selecting options in semi-MDPs and dynamically composing computation graphs. Lastly, we demonstrate the potential advantages of a society’s inherent modular structure for more efficient transfer learning.

You know that everything you think and do is thought and done by you. But what’s a “you”? What kinds of smaller entities cooperate inside your mind to do your work?

(Minsky, 1988)

¹Department of Computer Science, University of California, Berkeley, USA ²Department of Computer Science, Princeton University, USA. Correspondence to: Michael Chang <mbchang@berkeley.edu>.




Figure 1. We study how a society of primitive agents can be abstracted as a super-agent. The incentive mechanism is the abstraction barrier that relates the optimization problems of the super-agent with those of its constituent primitive agents.

1. Introduction

Biological processes, corporations, and ecosystems – physically decentralized, yet in some sense functionally unified. A corporation, for example, optimizes for maximizing profits as it were a single rational agent. But this agent abstraction is an illusion: the corporation is simply a collection of human agents, each solving their own optimization problems, most not even knowing the existence of many of their colleagues. But the human as the decision-making agent is also simply an abstraction of the trillions of cells making their own simpler decisions. *The society of agents is itself an agent*. What mechanisms bridge between these two levels of abstraction, and under what framework can we develop learning algorithms for studying the self-organizing nature of intelligent societies that pervade so much of the world?

Both the monolithic and the multi-agent optimization frameworks in machine learning offer a language for representing only one of the levels of abstraction but not the relation between both. The monolithic framework, the most commonly used in much of modern machine learning, considers a single agent that optimizes a single objective in an environment, whether it be minimizing classification loss or maximizing return. The multi-agent framework considers multiple agents that each optimize their own independent objective and each constitute each other’s learning environments. What distinguishes the multi-agent from the monolithic is the presence of multiple independent optimization problems. The difficulty of interpreting a learner in the monolithic framework as a society of simpler components is that all components are still globally coupled together by the same optimization problem without independent local

optimization problems themselves, as are the weights in a neural network trained by backpropagation. The difficulty of interpreting a multi-agent system under a global optimization problem is the computational difficulty of computing Nash equilibrium (Daskalakis et al., 2009), even for general two-player games (Chen et al., 2009).

To better understand the relationship between the society and the agent, this paper makes four contributions, each at a different level of abstraction. At the highest level, we define the **societal decision-making framework** to relate the local optimization problem of the agent to the global optimization problem of the society in the following restricted setting. Each agent is specialized to transform the environment from one state to another. The agents bid in an auction at each state and the auction winner transforms the state into another state, which it sells to the agents at the next time-step, thereby propagating a series of economic transactions. This framework allows us to ask what are properties of the auction mechanism and of the society that enable the global solution to a Markov decision process (MDP) that the society solves to emerge implicitly as a consequence of the agents optimizing their own independent auction utilities.

At the second level, we present a *solution* to this question by introducing the **cloned Vickrey society** that guarantees that the dominant strategy equilibrium of the agents coincides with the optimal policy of the society. We prove this result by leveraging the truthfulness property of the Vickrey auction (Vickrey, 1961) and showing that initializing redundant agents makes the primitives' economic transactions robust against market bubbles and suboptimal equilibria.

At the third level, we propose a class of **decentralized reinforcement learning algorithms** for optimizing the MDP objective of the society as an emergent consequence of the agents' optimizing their own auction utilities. These algorithms treat the auction utility as optimization objectives themselves, thereby learning a societal policy that is global in space and time using only credit assignment for learnable parameters that is local in space and time.

At the fourth level, we empirically investigate various *implementations* of the cloned Vickrey society under our decentralized reinforcement learning algorithm and find that a particular set of design choices, which we call the **credit conserving Vickrey** implementation, yields both the best performance at the societal and agent level.

Finally, we demonstrate that the societal decision making framework, along with its solution, the algorithm that learns the solution, and the implementation of this algorithm, is a broadly applicable perspective on self-organization to not only standard reinforcement learning but also selecting options in semi-MDPs (Sutton et al., 1999) and composing functions in dynamic computation graphs (Chang et al.,

2018). Moreover, we show evidence that the local credit assignment mechanisms of societal decision-making produce more efficient learning than the global credit assignment mechanisms of the monolithic framework.

2. Related Work

Describing an intelligent system as the product of interactions among many individual agents dates as far back as the Republic (Plato, 380 B.C.), in which Plato analyzes the human mind via an analogy to a political state. This theme continued into the early foundations of AI in the 1980s and 1990s through cognitive models such as the Society of Mind (Minsky, 1988) and Braitenberg vehicles (Braitenberg, 1986) and engineering successes in robotics (Brooks, 1991) and in visual pattern recognition (Selfridge, 1988).

The closest works to ours were the algorithms developed around that same time period that sought as we do to leverage a multi-agent society for achieving a global objective, starting as early as the bucket brigade algorithm (Holland, 1985), in which agents bid in a first-price auction to operate on the state and auction winners directly paid their bid to the winners from the previous step. Prototypical self-referential learning mechanisms (Schmidhuber, 1987) improved the bucket brigade by imposing credit conservation in the economic transactions. The neural bucket brigade (Schmidhuber, 1989) adapted the bucket brigade to learning neural network weights, where payoffs corresponded to weight changes. Baum (1996) observed that the optimal choice for an agent's bid should be equivalent to the optimal Q-value for executing that agent's transformation and developed the Hayek architecture for introducing new agents and removing agents that have gone broke. Kwee et al. (2001) added external memory to the Hayek architecture.

However, to this date there has been no proof to the best of our knowledge that the bid-updating schemes proposed in these works simultaneously optimize a global objective of the society in a decision-making context. Sutton (1988) provides a convergence proof for temporal difference methods that share some properties with the bucket brigade credit assignment scheme, but importantly does not take the competition between the individual agents into account. But it is precisely the competition among agents in multi-agent learning that make their equilibria nontrivial to characterize (Mazumdar et al., 2019). Our work offers an alternative auction mechanism for which we prove that the optimal solution for the global objective *does* coincide with a Nash equilibrium of the society. We follow similar motivations to Balduzzi (2014), which investigates incentive mechanisms for training a society of rational discrete-valued neurons. In contrast to other works that decouple the computation graph (Srivastava et al., 2013; Goyal et al., 2019; Peng et al., 2019; Pathak et al., 2019) but optimize a global ob-

jective, our work considers optimizing local objectives only. We consider economic transactions between time-steps, as opposed to within a single time-step (Ohsawa et al., 2018).

3. Preliminaries

To set up a framework for societal decision-making, we relate Markov decision processes (MDP) and auctions under a unifying language. We define an **environment** as a tuple that specifies an input space, an output space, and additional parameters for specifying an objective. An **agent** is a function that maps the input space to the output space. An **objective** is a functional that maps the learner to a real number. Given an environment and objective, the **problem** the agent solves is to maximize the value of the objective.

In the MDP environment, the input space is the state space \mathcal{S} and the output space is the action space \mathcal{A} . The agent is a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$. The transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, the reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and discount factor γ are additional parameters that specify the objective: the return $J(\pi) = \mathbb{E}_{\tau \sim p^\pi(\tau)} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right]$, where $p^\pi(\tau) = p(s_0) \prod_{t=0}^T \pi(a_t | s_t) \prod_{t=0}^{T-1} \mathcal{T}(s_{t+1} | s_t, a_t)$. The agent solves the problem of finding $\pi^* = \arg \max_\pi J(\pi)$. For any state s , the optimal action for maximizing $J(\pi)$ is $\pi^*(s) = \arg \max_a Q^*(s, a)$, where the optimal Q function $Q^*(s, a)$ is recursively defined as $Q^*(s, a) = \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [r(s, a) + \gamma \max_{a'} Q^*(s', a') | s, a]$.

In the auction environments we consider, the input space is a single auction item s and the output space is the bidding space \mathcal{B} . Instead of a single agent, each of N agents $\psi^{1:N}$ compete to bid for the auction item via its bidding policy $\psi^i : \{s\} \rightarrow \mathcal{B}$. Let \mathbf{b} be the vector of bids produced by $\psi^{1:N}$. The vector \mathbf{v}_s of each agent’s valuations for auction item s and the **action mechanism** – allocation rule $X : \mathcal{B}^N \rightarrow [0, 1]^N$ and pricing rule $P : \mathcal{B}^N \rightarrow \mathbb{R}_{\geq 0}^N$ – are additional parameters that specify each agent’s objective: the utility $U_s^i(\psi^{1:N}) = \mathbf{v}_s^i \cdot X^i(\mathbf{b}) - P^i(\mathbf{b})$, where $X^i(\mathbf{b})$ is the proportion of s allocated to i , and $P^i(\mathbf{b})$ is the scalar price i pays. Each agent i independently solves the problem of finding $\psi^{i*} = \arg \max_{\psi^i} U_s^i(\psi^{1:N})$. The independent optimization of objectives distinguishes a multi-agent problem from a single-agent one and makes multi-agent problems generally difficult to analyze when an agent’s optimal policy depends on the strategies of other agents.

However, if an auction is dominant strategy incentive compatible (DSIC), bidding one’s own valuation is optimal, independent of other players’ bidding strategies. That is, truthful bidding is the unique dominant strategy. Notably, the **Vickrey auction** (Vickrey, 1961), which sets $P^i(\mathbf{b})$ to be the second highest bid $\max_{j \neq i} \mathbf{b}^j$ and $X^i(\mathbf{b}) = 1$ if i wins and 0 and 0 respectively if i loses, is DSIC, which means the dominant strategy equilibrium occurs when every agent bids

truthfully, making the Vickrey auction straightforward to analyze. Another attractive property of the Vickrey auction is that the dominant strategy equilibrium automatically maximizes the social welfare $\sum_{i=1}^N \mathbf{v}^i \cdot X^i(\mathbf{b})$ (Roughgarden, 2016), which selects the bidder with the highest valuation as winner. The existence of dominant strategies in the Vickrey auction removes the need for agents to recursively model others, giving the Vickrey auction the practical benefit of running in linear time (Roughgarden, 2016).

4. Societal Decision-Making

The perspective of this paper is that a society of agents can be abstracted as an agent that itself solves an optimization problem at a global level as an emergent consequence of the optimization problems its constituent agents solve at the local level. To make this abstraction precise, we now introduce the **societal decision-making** framework for analyzing and developing algorithms that relate the global decision problem of a society to the local decision problems of its constituent agents. We use **primitive** and **society** to distinguish between the agents at the local and global levels, respectively, which we define in the context of their local and global environments and objectives:

Definition 4.1. A *primitive* ω is a tuple $(\psi, \phi_\mathcal{T})$ of a bidding policy $\psi : \mathcal{S} \rightarrow \mathcal{B}$ and transformation $\phi_\mathcal{T} : \mathcal{S} \rightarrow \mathcal{S}$.

Definition 4.2. A *society* Ω is a set of primitives $\omega^{1:N}$.

The **global environment** is an MDP that we call the **global MDP**, with state space \mathcal{S} and discrete action space $\mathcal{A} = \{1, \dots, N\}$ that indexes the primitives $\omega^{1:N}$. The **local environment** is an auction that we call the **local auction** with auction item $s \in \mathcal{S}$ and bidding space $\mathcal{B} = [0, \infty)$.

The connection between the local and global environments is as follows. Each state in the global MDP is an auction item for a different local auction. The winning primitive $\hat{\omega}$ of the auction at state s transforms s into the next state s' of the global MDP using its transformation $\phi_\mathcal{T}$, parameterized by the global MDP’s transition function \mathcal{T} . For each primitive i at each state s , its **local objective** is the utility $U_s^i(\psi^{1:N})$. Its **local problem** is to maximize $U_s^i(\psi^{1:N})$. The **global objective** is the return $J(\pi_\Omega)$ in the global MDP of the **global policy** π_Ω . The **global problem** for the society is to maximize $J(\pi_\Omega)$. We define the optimal **societal Q function** $Q_\Omega^*(s, \omega)$ as the expected return received from ω invoking its transformation $\phi_\mathcal{T}$ on s and the society activating primitives optimally with respect to $J(\pi_\Omega)$ afterward.

Since all decisions made at the societal level are an emergent consequence of decisions made at the primitive level, the societal decision-making framework is a self-organization perspective on a broad range of sequential decision problems. If each transformation $\phi_\mathcal{T}$ specifies a literal action, then societal decision-making is a decentralized re-framing

of standard reinforcement learning (RL). Societal decision-making also encompasses the decision problem of choosing $\phi_{\mathcal{T}s}$ as options in semi-MDPs (Sutton et al., 1999) as well as choosing $\phi_{\mathcal{T}s}$ as functions in a computation graph (Chang et al., 2018; Rosenbaum et al., 2017; Alet et al., 2018).

We are interested in auction mechanisms and learning algorithms for optimizing the global objective as an emergent consequence of optimizing the local objectives. Translating problems from one level of abstraction to another would provide a recipe for engineering a multi-agent system to achieve a desired global outcome and permit theoretical expectations on the nature of the equilibrium of the society, while giving us free choice on the architectures and learning algorithms of the primitive agents. To this end, we next present an auction mechanism for which the dominant strategy equilibrium of the primitives coincides with the optimal policy of the society, which we develop into a class of decentralized RL algorithms in later sections.

5. Mechanism Design for the Society

We first observe that to produce the optimal global policy, the optimal bidding strategy for each primitive at each local auction must be to bid their societal Q-value. By defining each primitive’s valuation of a state as its optimal societal Q-value at that state, we show that the Vickrey auction ensures the dominant strategy equilibrium profile of the primitives coincides with the optimal global policy. Then we show that a market economy perspective on societal decision-making overcomes the need to assume knowledge of optimal Q-values, although weakens the dominant strategy equilibrium to a Nash equilibrium. Lastly, we explain that adding redundant primitives to the society mitigates market bubbles by enforcing credit conservation. Proofs are in the Appendix.

5.1. Optimal Bidding

We state what was observed informally in (Baum, 1996):

Proposition 5.1. *Assume at each state s the local auction allocates $X^i(\mathbf{b}) = 1$ if i wins and $X^i(\mathbf{b}) = 0$ if i loses. Then all primitives ω^i bidding their optimal societal Q-values $Q_{\Omega}^*(s, \omega^i)$ collectively induce an optimal global policy.*

This proposition makes the problem of self-organization concrete: getting the optimal behavior in the global MDP to emerge from the optimal behavior in the local auctions can be reduced to incentivizing the primitives to bid their optimal societal Q-value at every state.

5.2. Dominant Strategies for Optimal Bidding

To incentivize the primitives to bid optimally, we propose to define the primitives’ valuations \mathbf{v}_s^i for each state s as their optimal societal Q-values $Q_{\Omega}^*(s, \omega^i)$ and use the Vickrey

auction mechanism for each local auction.

Theorem 5.2. *If the valuations \mathbf{v}_s^i for each state s are the optimal societal Q-values $Q_{\Omega}^*(s, \omega^i)$, then the society’s optimal global policy coincides with the primitives’ unique dominant strategy equilibrium under the Vickrey mechanism.*

Then, the utility $U_s^i(\psi^{1:N})$ at each state s that induces the optimal global policy, which we refer equivalently as $\hat{U}_s^i(\omega^{1:N})$ for the winning primitive $\hat{\omega}^i$ and $U_s^j(\omega^{1:N})$ for losing primitives ω^j , is given by

$$\hat{U}_s^i(\omega^{1:N}) = Q_{\Omega}^*(s, \hat{\omega}^i) - \max_{j \neq i} \mathbf{b}_s^j \quad (1)$$

and by $U_s^j(\omega^{1:N}) = 0$ for losing primitives.

5.3. Economic Transactions for Propagating Q_{Ω}^*

We have so far defined optimal bidding with respect to societal decision-making and characterized the utilities as functions of Q_{Ω}^* for which such bidding is a dominant strategy. We now propose to redefine the utilities without knowledge of Q_{Ω}^* by viewing the society as a market economy.

Monolithic frameworks for solving MDPs, such as directly optimizing the policy $J(\pi)$ with policy gradient methods, are analogous to **command-economies**, in which all production – the transformation of past states s_t into future states s_{t+1} – and wealth distribution – the credit assignment of reward signals to parameters – derive directly from single central authority – the MDP objective. In contrast, under the societal decision-making framework, the optimal global policy does *not* derive directly from the MDP objective, but rather emerges implicitly as the equilibrium of the primitives optimizing their own local objectives. We thus redefine the valuations \mathbf{v}_s^i following the analogy of a **market economy**, in which production and wealth distribution are governed by the economic transactions between the primitives.

Specifically, we couple the local auctions at consecutive time-steps in the same game by defining the valuation $\mathbf{v}_{s_t}^i$ of primitive $\hat{\omega}^i$ for winning the auction item s_t as the revenue it can receive in the auction at the next time-step by selling the product s_{t+1} of executing its transformation $\phi_{\mathcal{T}}^i$ on s_t . This compensation comes as the environment reward plus the (discounted) winning bid at the next time-step:

$$\underbrace{\hat{U}_{s_t}^i(\omega^{1:N})}_{\text{utility}} = r(s_t, \hat{\omega}^i) + \underbrace{\gamma \cdot \max_k \mathbf{b}_{s_{t+1}}^k}_{\text{revenue, or valuation } \mathbf{v}_{s_t}^i} - \underbrace{\max_{j \neq i} \mathbf{b}_{s_t}^j}_{\text{price}} \quad (2)$$

Analogous to a market economy, the revenue $\hat{\omega}^i$ receives for producing s_{t+1} from s_t depends on the price the winning primitive $\hat{\omega}^k$ at $t + 1$ is willing to bid for s_{t+1} . In turn, $\hat{\omega}^k$ sells s_{t+2} to the winning primitive at $t + 2$, and so on. Ultimately currency is grounded in the reward. Wealth is distributed based on what future primitives decide to bid for

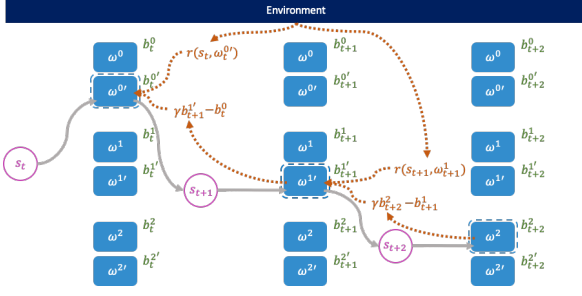


Figure 2. The cloned Vickrey society. In this market economy of primitive agents, wealth is distributed not directly from the global MDP objective but based on what future primitives decided to bid for the fruits of the labor of information processing carried out by past primitives transforming one state to another. The primitive $\omega_t^{0'}$ that wins the auction at time t receives an environment reward $r(s_t, \omega_t^{0'})$ as well as payment $b_{t+1}^{1'}$ from $\omega_{t+1}^{1'}$ for transforming s_t to s_{t+1} . By the Vickrey auction, the price $\omega_{t+1}^{1'}$ pays to transform s_{t+1} is the second highest bid $b_{t+1}^{1'}$ at time $t+1$. Because each primitive ω^i and its clone $\omega^{i'}$ have the same valuations, their bids are equivalent and so credit is conserved through time.

the fruits of the labor of information processing carried out by past primitives transforming one state to another.

Definition 5.1. A *Market MDP* is a global MDP in which all utilities are defined as in Equation 2.

As valuations now depend on the strategies of future primitives, the dominant strategy equilibrium from Theorem 5.2 must be weakened to a Nash equilibrium in the general case:

Proposition 5.3. In a *Market MDP*, it is a Nash equilibrium for every primitive to bid $Q_{\Omega}^*(s, \omega^i)$. Moreover, if the *Market MDP* is finite horizon, then bidding $Q_{\Omega}^*(s, \omega^i)$ is the unique Nash equilibrium that survives iterated deletion of weakly dominated strategies.

5.4. Redundancy for Credit Conservation

In general, credit is not conserved in the Market MDP: the winning primitive at time $t-1$ gets paid an amount equal to the highest bid at time t , but the winner at time t only pays an amount equal to the second highest bid at time t . At time t , ω^i could bid arbitrarily higher than $\max_{j \neq i} b^j$ without penalty, which distorts the valuations for primitives that bid before time t , creating a market bubble.

To prevent market bubbles, we propose a modification to the society, which we will call a **cloned society**, for enforcing credit conservation – for all transformations $\phi_{\mathcal{T}}$ initialize at least two primitives that share the same $\phi_{\mathcal{T}}$:

Lemma 5.4. For a cloned society, at the Nash equilibrium specified in Proposition 5.3, what the winning primitive $\hat{\omega}^i$ at time t receives from the winning primitive $\hat{\omega}^k$ at $t+1$ is exactly what $\hat{\omega}^k$ pays: $b_{s_{t+1}}^k$.

We now state our main result.

Theorem 5.5. Define a *cloned Vickrey society* as a cloned society that solves a Market MDP. Then it is a Nash equilibrium for every primitive in the cloned Vickrey society to bid $Q_{\Omega}^*(s, \omega^i)$. In addition, the price that the winning primitive pays for winning is equivalent to what it bid.

The significance of Theorem 5.5 is that guaranteeing truthful bidding of societal Q-values decouples the analysis of the local problem within each time-step from that of the global problem across time-steps, which opens the possibility for designing learners to reach a Nash equilibrium that we know exists. Without such a separation of these two levels of abstraction the entire society must be analyzed as a repeated game through time – a non-trivial challenge.

6. From Equilibria to Learning Objectives

So far the discussion has centered around the quality of the equilibria, which assumes the primitives know their own valuations. We now propose a class of decentralized RL algorithms for learning optimal bidding behavior without assuming knowledge of the primitives' valuations.

Instead of learning to optimize the MDP objective directly, we propose to train the primitives' parameterized bidding policies to optimize their utilities in Equation 2, yielding a class of decentralized RL algorithms for optimizing the global RL objective that is agnostic to the choice of RL algorithm used to train each primitive. By Theorem 5.5, truthful bidding for all agents is one global optimum to all agent's local learning problems that also serves as a global optimum for the society as whole. In the special case where the transformation $\phi_{\mathcal{T}}$ is a literal action, this class of decentralized RL algorithms can serve as an alternative to any standard algorithm for solving discrete-action MDPs. An on-policy learning algorithm is presented in Appendix D.

6.1. Local Credit Assignment in Space and Time

The global problem requires a solution that is global in space, because the society must collectively work together, and global in time, because the society must maximize expected return across time-steps. But an interesting property of using the auction utility in Equation 2 as an RL objective is that given redundant primitives it takes the form of the Bellman equation, thereby implicitly coupling all primitives together in space and time. Thus each primitive need only optimize for its immediate utility at a each time-step without needing to optimize for its own future utilities. This class of decentralized RL algorithms thus implicitly finds a global solution in space and time using only credit assignment that is *local in space*, because transactions are only between individual agents, and *local in time*, because each primitive need only solve a contextual bandit problem at each time-step.

6.2. Redundancy for Avoiding Suboptimal Equilibria

A benefit of casting local auction utilities as RL objectives is the practical development of learning algorithms that need not assume oracle knowledge of valuations, but unless care is taken with each primitives’ learning environments, the society may not converge to the globally optimal Nash equilibrium described in Section 5. As an example of a suboptimal equilibrium, in a Market MDP with two primitives, even if $v^1 = 1, v^2 = 2$, it is a Nash equilibrium for $b^1 = 100, b^2 = 0$. Without sufficient competitive pressure to not overbid or underbid, a rogue winner lacks the risk of losing to other primitives with similarly close valuations. Fortunately, the redundancy of a cloned Vickrey society serves a dual purpose of not only preventing market bubbles but also introducing competitive pressure in the primitives’ learning environments in the form of other clones.

7. Experiments

Now we study how well the cloned Vickrey society can recover the optimal societal Q-function as its Nash equilibrium. We compare several implementations of the cloned Vickrey society against baselines across simple tabular environments in Section 7.1, where the transformations $\phi_{\mathcal{T}}$ are literal actions. Then in Section 7.2 we demonstrate the broad applicability of the cloned Vickrey society for learning to select options in semi-MDPs and composing functions in a computation graph. Moreover, we show evidence for the advantages of the societal decision making framework over its monolithic counterpart in transferring to new tasks.

Although our characterization of the equilibrium and the learning algorithm are agnostic to the implementation of the primitive, in our experiments the bidding policy ψ is implemented as a neural network that maps the state to the parameters of a Beta distribution, from which a bid is sampled. We used proximal policy optimization (PPO) (Schulman et al., 2017) to optimize the bidding policy parameters.

7.1. Numerical Simulations

In the following simulations we ask: (1) How closely do the bids the primitives learn match their optimal societal Q-values? (2) Does the solution to the global objective emerge from the competition among the primitives? (3) How does redundancy affect the solutions the primitives converge to?

We first consider the *Market Bandit*, the simplest global MDP with one state, and compare the cloned Vickrey society with societies with different auction mechanisms. Next we consider the *Chain* environment, a sparse-reward multi-step global MDP designed to study market bubbles. Last we consider the *Duality* environment, a multi-step global MDP designed to study suboptimal equilibria. We use *solitary society* to refer to a society without redundant clones.

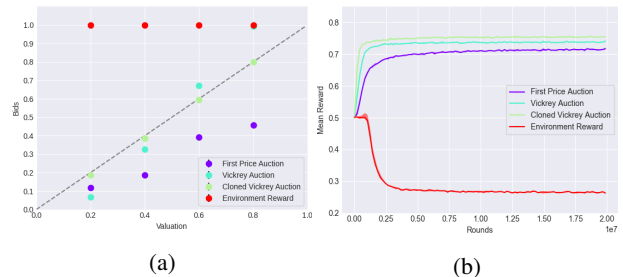



Figure 3. **Market Bandits.** We compare the cloned Vickrey society (Cloned Vickrey Auction) against solitary societies that use the first-price auction mechanism (first price auction), the Vickrey auction mechanism (Vickrey Auction), and a mechanism whose utility is only the environment reward (Environment Reward). The dashed line in (a) indicates truthful bidding of valuations. The cloned Vickrey society’s bids are closest to the true valuations which also translates into the best global policy (b).

7.1.1. MARKET BANDITS AND AUCTION MECHANISMS

The simulation primarily studies question (1) by comparing the auction mechanism of the cloned Vickrey society against other mechanisms in eliciting bids that match optimal societal Q-values. We compare against *first price auction*, based on Holland (1985), a solitary society that uses the first price auction mechanism for the local auction, in which the winning primitive pays a price of their bid, rather than the second highest bid; against *Vickrey Auction*, a solitary version of the cloned Vickrey society; and against *Environment Reward*, a baseline solitary society whose utility function uses only the environment reward, with no price term $P^2(\mathbf{b})$.

The environment is a four-armed Market Bandit whose arms correspond to transformations ϕ that deterministically yield reward values of 0.2, 0.4, 0.6, and 0.8. Solitary societies thus would have four primitives, while cloned societies have eight. Since we are mainly concerned with question (1), we stochastically drop out a subset of primitives at every round to give each primitive a chance to win and learn its value.

In the Market Bandit, the arm rewards directly specify the primitives’ valuations, so if the primitives successfully learned their valuations, we would expect them to learn to bid values exactly at 0.2, 0.4, 0.6, and 0.8. Figure 3a shows that the primitives in the cloned Vickrey society learn to bid most closely to their true valuations, which also translates into the best global policy in Figure 3b, answering question (2). To address question (3), we observe that the solitary Vickrey society’s bids are more spread out than those of the cloned Vickrey society because there is no competitive pressure to learn the valuations exactly. This is not detrimental when the global MDP has only one time-step, but the next section shows that the lack of redundancy creates market bubbles that result in a suboptimal global policy.



	$\hat{\omega}_t$ Receives	$\hat{\omega}_{t+1}$ Pays
CCV	\mathbf{b}'_{t+1}	\mathbf{b}'_{t+1}
BB	$\hat{\mathbf{b}}_{t+1}$	$\hat{\mathbf{b}}_{t+1}$
V	$\hat{\mathbf{b}}_{t+1}$	\mathbf{b}'_{t+1}

Figure 4. Implementations. The table shows the bid price that temporally consecutive winners $\hat{\omega}_{t+1}$ and $\hat{\omega}_t$ pay and receive based on three possible implementations of the cloned Vickrey society: *CCV*, *BB*, *V*, with tradeoffs depicted in the Venn diagram. We use $\hat{\mathbf{b}}_{t+1}$ and \mathbf{b}'_{t+1} to denote the highest and second highest bids at time $t + 1$ respectively.

7.1.2. MARKET MDPs

Now we consider MDPs that involve multiple time-steps to understand how learning is affected when primitives' valuations are defined by the bids of future primitives. Redundancy theoretically makes what the winner $\hat{\omega}_t$ receives and $\hat{\omega}_{t+1}$ equivalent, but the stochasticity of the bid distribution yields various possible implementations for the cloned Vickrey society in the market MDP—*bucket brigade* (*BB*), *Vickrey* (*V*), and *credit conserving Vickrey* (*CCV*), summarized in Figure 4, with different pros and cons. Note that the solitary bucket brigade society is a multiple time-step analog of the *first price auction* in Section 7.1.1. We aim to empirically test which of the implementations of the cloned Vickrey society yields best local and global performance.

The Chain environment and market bubbles. A primary purpose of the *Chain* environment (Figure 6a) is to study the effect of redundancy on mitigating market bubbles in the bidding behavior and how that affects global optimality. The *Chain* environment is a sparse reward finite-horizon environment which initializes the society at s_0 on the left, yields a terminal reward of 0.8 if the society enters the goal state s_5 on the right, and ends the episode after 20 steps if the society does not reach the goal. *Chain* thus tests the ability of the society to propagate utility from future agents to past agents in the absence of an immediate environment reward signal. We compare the bidding behaviors of the *BB*, *V*, and *CCV* implementations of the cloned Vickrey society, as well as those of their solitary counterparts, in Figure 6 and the society's global learning curve in Figure 8a.

To answer question (1), we observe that redundancy indeed prevents market bubbles, with the cloned *CCV* implementation bidding closest to the optimal societal Q values. Details are in the caption of Figure 5. When we consider the society's global learning curve in Figure 8a, the answers to questions (2) and (3) go hand-in-hand: the solitary societies fail to find the globally optimal policy and the cloned *CCV* implementation has the highest sample efficiency.

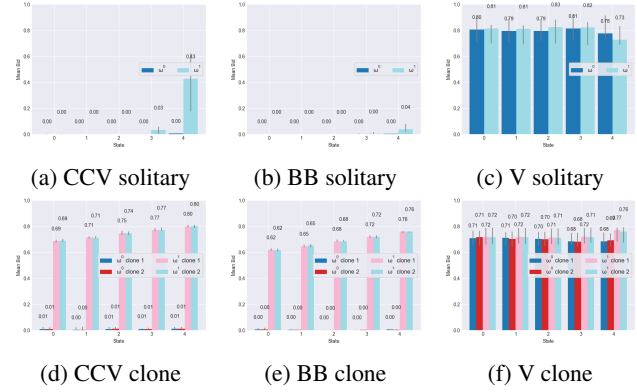


Figure 5. Learned Bidding Strategies for Chain. We organize the analysis by distinguishing between the credit-conserving (*CCV* and *BB*) and the non-credit-conserving (*V*) implementations. The solitary *CCV* (a) and *BB* (b) implementations learn to bid very close to 0: *CCV* because the valuation for a primitive at t is only the second-highest bid at $t + 1$, resulting in a rapid decay in the valuations leftwards down the chain; *BB* because each primitive is incentivized to pay as low of a price for winning as possible. The cloned *CCV* (d) and *BB* (e) implementations learn to implement a form of return decomposition (Arjona-Medina et al., 2019) that redistributes the terminal reward into a series of positive payoffs back through the chain, each agent getting paid for contributing to moving the society closer to the goal state, where the *CCV* implementation's bids are closer to the optimal societal Q-value than those of the *BB* implementation. Because both the solitary (c) and cloned (f) versions of the *V* implementations do not conserve credit, they learn to bid close to the optimal societal Q-value, but both suffer from market bubbles where the primitive for going left bids higher than the primitive for going right, even though the optimal global policy is to keep moving right.

The Duality environment and suboptimal equilibria. The *Chain* environment experiments suggested a connection between lack of redundancy and globally suboptimal equilibria, the subtleties of which we explore further in the *Duality* environment (Figure 6b). The *CCV* implementation has yielded the best performance so far but does not guarantee Bellman optimality (Figure 4) without redundant primitives. We show that in the *Duality* environment, without redundant primitives the dominant strategy equilibrium would lead the society to get stuck in a self-loop at s_1 indefinitely, even though the global optimal solution would be to cycle back and forth between s_0 and s_1 . The reasoning for this is explained in Appendix F.1.2. Figure 8b indeed shows that a society with redundant primitives learns a better equilibrium than without.

7.2. Semi-MDPs and Computation Graphs

One benefit of framing global decision-making from the perspective of local economic transactions is that the same societal decision-making framework and learning algorithms




Figure 6. Multi-Step MDPs. (a) In the *Chain* environment, the society starts at state s_0 and the goal state is s_5 . Only activating primitive ω_1 at state s_4 yields reward. The optimal global policy is to directly move right by continually activating ω_1 . Without credit conservation, the society may get stuck going back and forth between s_0 and s_4 without reaching the goal. (b) In the *Duality* environment, the society starts at state s_0 . s_{-1} is an absorbing state with perpetual negative rewards. The optimal societal policy is to cycle between s_0 and s_1 to receive unbounded reward, but without redundant primitives, the society may end up in a suboptimal perpetual self-loop at s_1 .

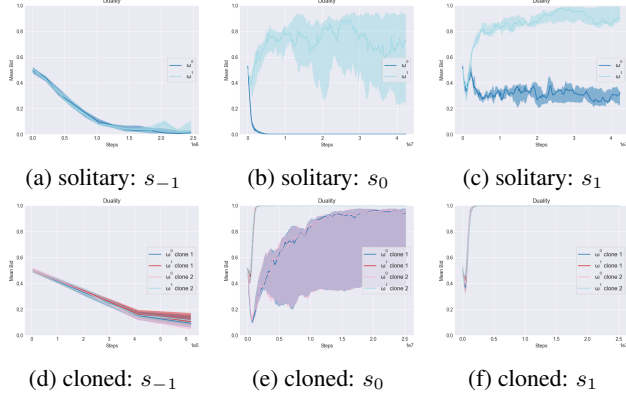


Figure 7. CCV Bidding Curves for *Duality*. Each column shows the bidding curves of the solitary (top row) and cloned (bottom row) CCV societies for states s_{-1} , s_0 , and s_1 . Without redundant primitives to force the second-highest and highest valuations to be equal, the dominant strategy of truthful bidding may not coincide with the globally optimal policy because the solitary CCV implementation does not guarantee Bellman optimality. The bidding curves in (c) show that ω^1 learns a best response of bidding *higher* than primitive ω^0 at state s_1 , even though it would be globally optimal for the society if ω^0 wins at s_1 . Adding redundant primitives causes the second-highest and highest valuations to be equal, causing ω^0 to learn to bid highly as well at s_1 , which results in a more optimal return as shown in Figure 8b.

can be used regardless of the type of transformation $\phi_{\mathcal{T}}$. We now show in the *Two Rooms* environment that the cloned Vickrey society can learn more efficiently than a monolithic counterpart to select among pre-trained options to solve a *gym-minigrid* (Chevalier-Boisvert et al., 2018) navigation task that involves two rooms separated by a bottleneck state. We also show in the *Mental Rotation* environment that the cloned Vickrey society can learn to dynamically compose computation graphs of pre-specified affine transformations for classifying spatially transformed MNIST digits. We use the cloned CCV society for these experiments.

7.2.1. TRANSFERRING WITH OPTIONS IN SEMI-MDPs

We construct a two room environment, *Two Rooms*, which requires opening the red door and reaching either a green

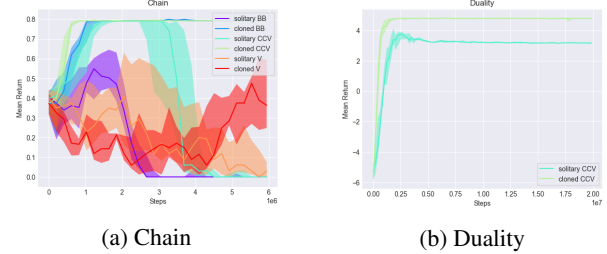


Figure 8. Multi-Step MDP Global Learning Curves. We observe that cloned societies are more robust against suboptimal equilibria than solitary societies. Furthermore the cloned CCV implementation achieves the best sample efficiency, suggesting that truthful bidding and credit-conservation are important properties to enforce for enabling the optimal global policy to emerge.

goal or blue goal. The transformations $\phi_{\mathcal{T}}$ are subpolicies that have been pre-trained to open the red door, reach the green goal, and reach the blue goal. In the pre-training task, only reaching the green goal gives a non-zero terminal reward and reaching the blue goal does not give reward. In the transfer task, the rewards for reaching the green and blue goals are switched. We compared the cloned Vickrey society against a non-hierarchical monolithic baseline that selects among the low-level *gym-minigrid* actions as well as a hierarchical monolithic baseline that selects among the pre-trained subpolicies. Both baseline policies sample from a Categorical distribution and are trained with PPO. The cloned Vickrey society is more sample efficient in learning on the pre-training task and significantly faster in adapting to the transfer task (Figure 10).

Our hypothesized explanation for this efficiency is that the local credit assignment mechanisms of a society parallelize learning across space and time, a property that is not true of the global credit assignment mechanisms of a monolithic learner. To test this hypothesis, we observe in the bottom-right of Figure 10 that not only has a higher percentage of the hierarchical monolithic baseline’s weights shifted during transfer compared to our method but they also shift to a larger degree, which suggests that the hierarchical monolithic baseline’s weights are more globally coupled

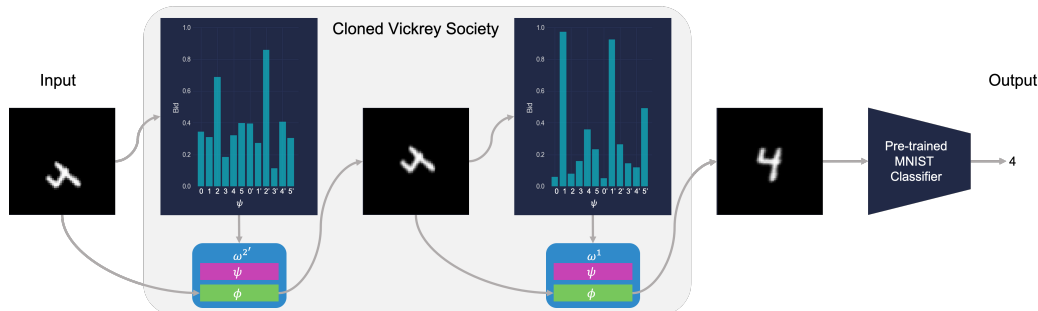


Figure 9. Mental Rotation. The cloned Vickrey society learns to transform the image into a form that can be classified correctly with a pre-trained classifier by composing two of six possible affine transformations of rotation and translation. Clones are indicated by an apostrophe. In this example, the society activated primitive $\omega^{2'}$ to translate the digit up then primitive ω^1 to rotate the digit clockwise. Though the bidding policies ψ^i and $\psi^{i'}$ of the clones ω^i and $\omega^{i'}$ have the same parameters, their sampled bids may be different because the bidding policies are stochastic.

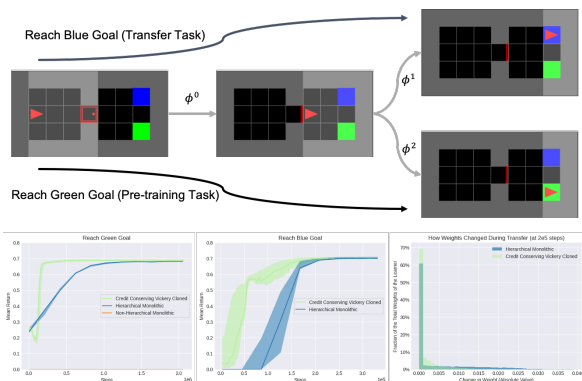


Figure 10. Two Rooms. The cloned Vickrey society adapts more quickly than the hierarchical monolithic baseline in both the pre-training and the transfer tasks. The bottom-right figure, which is a histogram of the absolute values of how much the weights have shifted from fine-tuning on the transfer task, shows that more weights shift, and to a larger degree, in the hierarchical monolithic baseline than in our method. This seems to suggest that the cloned Vickrey society is a more modular learner than the hierarchical monolithic baseline. The non-hierarchical monolithic baseline does not learn to solve the task from scratch.

and perhaps thereby slower to transfer. While this analysis is not comprehensive, it is a suggestive result that motivates future work for further study.

7.2.2. COMPOSING DYNAMIC COMPUTATION GRAPHS

We adapt the Image Transformations task from Chang et al. (2018) as what we call the *Mental Rotation* environment (Figure 9), in which MNIST images have been transformed with a composition of one of two rotations and one of four translations. There are $60,000 \times 8 = 240,000$ possible unique inputs, meaning learning an optimal global policy involves training primitives across 240,000 local auctions. The transformations $\phi_{\mathcal{T}}$ are pre-specified affine transforma-

tions. The society must transform the image into a form that can be classified correctly with a pre-trained classifier, with a terminal reward of 1 if the predicted label is correct and 0 otherwise. The cloned Vickrey society converges to a mean return of 0.933 with a standard deviation of 0.014.

8. Conclusion

This work formally defines the societal decision-making framework and proves the optimality of a society – the cloned Vickrey society – for solving a problem that was first posed in the AI literature in the 1980s: to specify the incentive structure that causes the global solution of a society to emerge as the equilibrium strategy profile of self-interested agents. For training the society, we further proposed a class of decentralized reinforcement learning algorithms whose global objective decouples in space and time into simpler local objectives. We have demonstrated the generality of this framework for selecting actions, options, and computations as well as its potential advantages for transfer learning.

The generality of the societal decision-making framework opens much opportunity for future work in decentralized reinforcement learning. A society’s inherent modular structure suggests the potential of reformulating problems with global credit-assignment paths into problems with much more local credit-assignment paths that offer more parallelism and independence in the training of different components of a learner. Understanding the learning dynamics of multi-agent societies continues to be an open problem of research. It would be exciting to explore algorithms for constructing and learning societies in which the primitives are also societies themselves. We hope that the societal decision-making framework and its associated decentralized reinforcement learning algorithms provide a foundation for future work exploring the potential of creating AI systems that reflect the collective intelligence of multi-agent societies.

Acknowledgements

The authors are grateful to the anonymous reviewers, Michael Janner, and Glen Berseth for their feedback on this paper, as well as Aviral Kumar, Sam Toyer, Anirudh Goyal, Marco Cusumano-Towner, Abhishek Gupta, and Dinesh Jayaraman for helpful discussions. This research was supported by AFOSR grant FA9550-18-1-0077, NSF CCF-1717899, and Google Cloud Platform. MC is supported by the National Science Foundation Graduate Research Fellowship Program. SK is supported by Berkeley Engineering Student Services.

References

- Alet, F., Lozano-Pérez, T., and Kaelbling, L. P. Modular meta-learning. *arXiv preprint arXiv:1806.10166*, 2018.
- Arjona-Medina, J. A., Gillhofer, M., Widrich, M., Unterthiner, T., Brandstetter, J., and Hochreiter, S. Rudder: Return decomposition for delayed rewards. In *Advances in Neural Information Processing Systems*, pp. 13544–13555, 2019.
- Balduzzi, D. Cortical prediction markets. *arXiv preprint arXiv:1401.1465*, 2014.
- Baum, E. B. Toward a model of mind as a laissez-faire economy of idiots. In *ICML*, pp. 28–36, 1996.
- Bellmann, R. Dynamic programming princeton university press. *Princeton, NJ*, 1957.
- Braitenberg, V. *Vehicles: Experiments in synthetic psychology*. 1986.
- Brooks, R. A. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.
- Chang, M. B., Gupta, A., Levine, S., and Griffiths, T. L. Automatically composing representation transformations as a means for generalization. *arXiv preprint arXiv:1807.04640*, 2018.
- Chen, X., Deng, X., and Teng, S.-H. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.
- Chevalier-Boisvert, M., Willems, L., and Pal, S. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- Goyal, A., Sodhani, S., Binas, J., Peng, X. B., Levine, S., and Bengio, Y. Reinforcement learning with competitive ensembles of information-constrained primitives. *arXiv preprint arXiv:1906.10667*, 2019.
- Holland, J. H. Properties of the bucket brigade. In *Proceedings of an International Conference on Genetic Algorithms and their Applications 1-7*, 1985.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kwee, I., Hutter, M., and Schmidhuber, J. Market-based reinforcement learning in partially observable worlds. In *International Conference on Artificial Neural Networks*, pp. 865–873. Springer, 2001.
- Mazumdar, E., Ratliff, L. J., Jordan, M. I., and Sastry, S. S. Policy-gradient algorithms have no guarantees of convergence in continuous action and state multi-agent settings. *arXiv preprint arXiv:1907.03712*, 2019.
- Minsky, M. *Society of mind*. Simon and Schuster, 1988.
- Ohsawa, S., Akuzawa, K., Matsushima, T., Bezerra, G., Iwasawa, Y., Kajino, H., Takenaka, S., and Matsuo, Y. Neuron as an agent, 2018. URL <https://openreview.net/forum?id=BkfEzz-0->.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Pathak, D., Lu, C., Darrell, T., Isola, P., and Efros, A. A. Learning to control self-assembling morphologies: a study of generalization via modularity. In *Advances in Neural Information Processing Systems*, pp. 2295–2305, 2019.
- Peng, X. B., Chang, M., Zhang, G., Abbeel, P., and Levine, S. Mcp: Learning composable hierarchical control with multiplicative compositional policies. *arXiv preprint arXiv:1905.09808*, 2019.
- Plato. *Republic*. 380 B.C.
- Rosenbaum, C., Klinger, T., and Riemer, M. Routing networks: Adaptive selection of non-linear functions for multi-task learning. *arXiv preprint arXiv:1711.01239*, 2017.
- Roughgarden, T. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.
- Schmidhuber, J. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta... hook*. PhD thesis, Technische Universität München, 1987.
- Schmidhuber, J. A local learning algorithm for dynamic feedforward and recurrent networks. *Connection Science*, 1(4):403–412, 1989.

- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Selfridge, O. G. *Pandemonium: A Paradigm for Learning*, pp. 115122. MIT Press, Cambridge, MA, USA, 1988. ISBN 0262010976.
- Srivastava, R. K., Masci, J., Kazerounian, S., Gomez, F., and Schmidhuber, J. Compete to compute. In *Advances in neural information processing systems*, pp. 2310–2318, 2013.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Sutton, R. S., Precup, D., and Singh, S. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211, 1999.
- Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.