

Network Preference Dynamics using Lattice Theory

Hans Riess¹ Gregory Henselman-Petrusek² Michael C. Munger³ Robert Ghrist⁴
Zachary I. Bell⁵ Michael M. Zavlanos^{1,6}

Abstract—Preferences, fundamental in all forms of strategic behavior and collective decision-making, in their raw form, are an abstract ordering on a set of alternatives. Agents, we assume, revise their preferences as they gain more information about other agents. Exploiting the ordered algebraic structure of preferences, we introduce a message-passing algorithm for heterogeneous agents distributed over a network to update their preferences based on aggregations of the preferences of their neighbors in a graph. We demonstrate the existence of equilibrium points of the resulting global dynamical system of local preference updates and provide a sufficient condition for trajectories to converge to equilibria: stable preferences. Finally, we present numerical simulations demonstrating our preliminary results.

I. INTRODUCTION

When agents or teams of agents in multi-agent systems are responsible for unilateral decision making, a number of feasible options are ranked based on independent information sources as well as the biases of agents. When there is not enough information or time to adequately compare all feasible options, a decision dilemma ensues [1]. In this paper, we design a decentralized mechanism for multi-agent systems to locally share preferences in order to augment each agent’s independent decision-making capabilities and, when possible, avoid operational paralysis. As our main contribution, we introduce and offer a qualified solution to the *preference dynamics problem* which asks how preferences of agents evolve as they gather information and interact with other agents. This challenge is distinguished from the *social choice problem*, to design (centralized) preference-aggregation mechanisms that satisfy criteria such as fairness or resistance to strategic manipulation (for instance, see [2]).

In their traditional use in economics and social science, preferences are commonly associated with subjective taste. Here, we take the liberty to extend the notion of preference to arbitrary decision-making contexts, ranging from reinforcement learning (RL) [3] to games [4] to language models [5]. Formally, a *preference relation* is a relative ordering on a set of *alternatives* (i.e. possible choices or options). While preference relations represented by utility

functions¹(which is in itself a strong assumption since even complete preferences can be unrepresentable [6]) come with a slew of analytical techniques and computational algorithms, e.g. a Nash equilibrium, raw preference relations have no such exploitable analytical structure. To analyze raw preference relations, in this paper we resort to algebraic techniques, rather than analytic, posing a new set of challenges. However, our primary motivation for avoiding utility theory is that time-variant utility functions cannot capture preference formation since cardinal or ordinal utilities necessarily compare every alternative. Moreover, the distinction between indifference and indecisiveness, i.e., the agent has not made up its mind or faced a choice [7], is captured by incomplete preference relations and not utility functions, although extended preferences are another approach here [8].

We propose an entirely new methodology that relies on imposing an information order on preference relations (Section III) to reason about incomplete preferences which, also, is capable of modeling preference change (Section IV). Mechanisms of preference change include revision, contradiction, as well as addition and subtraction of alternatives [9]. Our message-passing model of preference dynamics, based on the algebraic lattice structure of preference relations, incorporates elements of both revision (the join operation) and contradictions (the meet operation). Avoiding the philosophical question as to whether preference can in fact change (see [10]), our preference dynamics model can be viewed as a process for how preferences are formed in the first place as new information is revealed. The “true” preferences of agents, then, are, perhaps, those that are stable: they are equilibrium points in a preference dynamical system. We characterize the stable preferences and discuss sufficient conditions for when they are computable (Section V).

Closely related to the preference dynamics problem considered here are recent efforts on consensus [11] and synchronization [12] using lattice orders. Other works have used lattice theory to model preferences [13], although the lattice structure they introduce is different from ours. Several authors have studied aggregation or consensus on lattices [14], [15], and there is a corpus of literature addressing the design of preference-aggregation mechanisms, originating with [16], but, to our knowledge, none of these works have approached the problem of preference dynamics from either a decentralized or algebraic point-of-view, both of which we do here. The several efforts to formalize classical consensus

¹Dept. ECE, Duke University, Durham, NC, USA.

²Pacific Northwest National Lab, Richland, WA, USA.

³Dept. Political Science, Duke University, Durham, NC, USA.

⁴Dept. ESE, University of Penn., Philadelphia, PA, USA.

⁵Air Force Research Laboratory, Eglin AFB, FL, USA.

⁶Dept. Mech. Eng. & Material Sci, Duke University, Durham, NC, USA.

Corresponding author email: hans.riess@duke.edu.

This work is supported in part by AFOSR under award #FA9550-19-1-0169 and by ONR under agreement #N00014-18-1-2374.

¹A preference relation \succsim is representable by a utility function u if $u(a) \geq u(b)$ if and only if $a \succsim b$.

algorithms with lattice theory [17], motivated by data science [18] and distributed systems [19], are restricted to interaction graphs with a complete or star topology.

Our treatment of preference dynamics is also closely related to opinion dynamics (see [20] for a survey), originating with [21]. Opinions are typically framed as numeric representations of likes or dislikes, versus preferences which depict comparisons. Several of the mechanisms for preference change we discuss (Section IV) have analogues in opinion dynamics including lying, exaggerating, or downplaying [22], stubbornness [23], incorporating a notion of confidence in other agents' opinions [24], [25], and confirmation bias [26].

II. PROBLEM DEFINITION

Suppose agents are collected in a finite set $\mathcal{N} = \{1, 2, \dots, N\}$ and interact with other agents through a fixed undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. The set $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ is the set of *edges* of the graph, $(i, j) \in \mathcal{E}$ if Agent i and Agent j interact. For $i \in \mathcal{N}$, let $\mathcal{N}_i = \{j : (i, j) \in \mathcal{E}\}$ denote the set of *neighbors*. Agents form preferences over a fixed set of alternatives, denoted by \mathcal{A} . The set of all feasible preference relations over a fixed set of alternatives \mathcal{A} is denoted by $\text{Pre}(\mathcal{A})$. Preference relations (elements) in the set $\text{Pre}(\mathcal{A})$ are denoted by $a \succsim_i b$ (instead of as an ordered pair) and indexed by each agent $i \in \mathcal{N}$. Thus, the expression $a \succsim_i b$ conveys the meaning that *Agent i prefers Alternative a to Alternative b* . We say that an agent is indifferent between two alternatives, written $a \sim_i b$, if both $a \succsim_i b$ and $b \succsim_i a$ hold. We assume that agents' preferences satisfy the following minimal axioms.

Assumption 1 (Reflexivity): The preference \succsim_i satisfies $a \succsim_i a$ for all $a \in \mathcal{A}$ and for all $i \in \mathcal{N}$.

Assumption 2 (Transitivity): For all $i \in \mathcal{N}$ and for all $a, b, c \in \mathcal{A}$ such that $a \succsim_i b$ and $b \succsim_i c$, it is necessarily true that also $a \succsim_i c$.

Irreflexive preferences are said to be *strict* while reflexive preferences are *weak*. One can always obtain a strict preference from a weak preference by defining $a \succ_i b$ if and only $a \succsim_i b$ and $b \not\succsim_i a$. Intransitive preferences are inconsistent. We recall from the introduction that we do *not* assume preferences are complete, i.e., we do not assume for all $a, b \in \mathcal{A}$, $a \succsim_i b$ or $b \succsim_i a$. Thus, an agent may be indecisive, perhaps making comparisons between alternatives with limited information.

A. Preference Dynamics

Let $\boldsymbol{\pi} = (\succsim_1, \succsim_2, \dots, \succsim_N)$ be a tuple of preference relations indexed by each agent, called a *preference profile*. We study how a preference profile changes based on interactions between agents. If $J \subseteq \mathcal{N}$ is a group of agents, we let π_J be the preference profile restricted the indices in J ; in particular, if $J = \{j\}$, π_j is simply \succsim_j . Agents update their preferences iteratively at discrete time instants $t = 0, 1, 2, \dots$ according to coupled dynamics

$$\begin{aligned} \pi_i(t+1) &= F_i(\boldsymbol{\pi}(t)) \\ \pi_i(0) &= \pi_{i,0} \end{aligned}, \quad i \in \mathcal{N} \quad (1)$$

where $F_i : \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A}) \rightarrow \text{Pre}(\mathcal{A})$.

The first problem we address in this paper is modeling preference dynamics over a graph. Unlike opinion dynamics, in which standard dynamical systems theory can model how likes and dislikes on topics change over time, preference dynamics requires some level of algebraic sophistication. Thus, notions of structure-preserving maps as well as binary operations on preference relations are introduced with no assumed background (Section III). At the same time, if we suppose that agents update their preferences according to the dynamics in (1), we also study the class of maps F_i that preserve the structure of preferences and model the process of preference revision. Specifically, in designing F_i , we focus on several aspects of preference change including agents having different "personalities" informing how they communicate, aggregate, and update preferences, agents fairly incorporating new comparisons into their preferences given a high enough level of consensus, all the while maintaining the consistency (transitivity) of their own preferences.

The second problem we address is a computational one. Specifically, assuming that preferences evolve according to the coupled dynamical system (1) and the iterative maps F_i are known, we analyze the existence of equilibrium points and discuss ways to compute them. Moreover, we provide sufficient conditions for the dynamics to converge to an equilibrium point. We interpret these equilibrium points as stable preference profiles under a model of preference dynamics.

III. PREFERENCE LATTICES

In this section, we show that the set of preferences on a fixed set of alternatives satisfying Assumption 1-2 is an algebraic structure known as a *lattice* (e.g. see [27]). We, then, interpret the meaning of the binary operations of the lattice, called *meet* and *join*, as two ways for agents to amalgamate preferences. But first, we present some rudimentary material about lattices, starting with even more general binary relations.

A. Preorders & Partial Orders

A *preorder* is a set \mathcal{P} with a relation \preceq satisfying the axioms of transitivity and reflexivity. Preferences, we assume, are preorders, and we will use both terms interchangeably throughout this paper. A *partial order* is a preorder, with the order relation usually written \preceq , satisfying a third axiom: if $\pi_1 \preceq \pi_2$ and $\pi_2 \preceq \pi_1$, then $\pi_1 = \pi_2$. Note, we do not assume preferences satisfy this third axiom, because agents can be indifferent between two alternatives. In partial orders, however, we conveniently write $\pi_1 \prec \pi_2$ whenever $\pi_1 \preceq \pi_2$, but $\pi_2 \not\preceq \pi_1$.

We consider several types of maps between ordered sets. A map $f : \mathcal{P} \rightarrow \mathcal{Q}$ between partial orders is *monotone* if $\pi_1 \preceq_{\mathcal{P}} \pi_2$ implies $f(\pi_1) \preceq_{\mathcal{Q}} f(\pi_2)$. A map $f : \mathcal{P} \rightarrow \mathcal{P}$ is *inflationary* if $f(\pi) \succeq \pi$ for all $\pi \in \mathcal{P}$, and *deflationary* if $f(\pi) \preceq \pi$ for all $\pi \in \mathcal{P}$. Note, inflationary or deflationary functions are not necessarily monotone. Cartesian products of partial orders $\mathcal{P}_1 \times \mathcal{P}_2 \times \dots \times \mathcal{P}_N$ yield a partial order called the *product order*: $(\pi_1, \pi_2, \dots, \pi_N) \preceq (\pi'_1, \pi'_2, \dots, \pi'_N)$ if and only if $\pi_i \preceq \pi'_i$ for all $i = 1, 2, \dots, N$.

B. Lattices

Lattices are partial orders with a rich algebraic structure given by two merging operations called “meet” and “join.”

Definition 1 (Lattices): A lattice is a partial order (\mathcal{L}, \preceq) such that, for any two elements $\pi_1, \pi_2 \in \mathcal{L}$, the operations

$$\begin{aligned}\pi_1 \wedge \pi_2 &= \max\{\pi \in \mathcal{L} : \pi \preceq \pi_1, \pi \preceq \pi_2\}, \\ \pi_1 \vee \pi_2 &= \min\{\pi \in \mathcal{L} : \pi \succeq \pi_1, \pi \succeq \pi_2\},\end{aligned}$$

called *meet* (greatest lower bound) and *join* (least upper bound), exist.

We recall a set of properties that characterize lattices as ordered algebraic structures.

Lemma 1: Suppose (\mathcal{L}, \preceq) is a lattice. Then, $\wedge, \vee : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ satisfy the following:

- (i) *commutativity*, i.e., $\pi_1 \wedge \pi_2 = \pi_2 \wedge \pi_1$, $\pi_1 \vee \pi_2 = \pi_2 \vee \pi_1$;
- (ii) *associativity*, i.e., $\pi_1 \wedge (\pi_2 \wedge \pi_3) = (\pi_1 \wedge \pi_2) \wedge \pi_3$, etc.;
- (iii) *idempotence*, i.e., $\pi \wedge \pi = \pi \vee \pi = \pi$;
- (iv) *absorption*, i.e., $\pi_1 \vee (\pi_1 \wedge \pi_2) = \pi_1 \wedge (\pi_1 \vee \pi_2) = \pi_1$;
- (v) *monotonicity*, i.e., \vee and \wedge are monotone in both arguments.

Proof: See [28, §I.5 Lemma 1, §I.5 Lemma 3]. ■

For lattices that are not finite, there is a distinction between the existence of binary meets and joins versus arbitrary meets and joins. Given $\{x_j\}_{j \in J} \subseteq \mathcal{L}$, a lattice is *complete* if $\bigwedge_{j \in J} \pi_j = \max\{\pi \in \mathcal{L} : \pi \preceq \pi_j \forall j \in J\}$ and $\bigvee_{j \in J} \pi_j = \min\{\pi \in \mathcal{L} : \pi \succeq \pi_j \forall j \in J\}$ exist. In particular, $\bigwedge \emptyset = \top$, the *maximum element*, and $\bigvee \emptyset = \perp$, the *minimum element*. By the obvious induction argument, finite lattices are complete.

C. The Information Order

As discussed in Section II, let $\text{Pre}(\mathcal{A})$ denote the set of preorders over a ground set \mathcal{A} , called *preference relations*. We equip $\text{Pre}(\mathcal{A})$ with the partial order inherited from the inclusion order on the powerset of $\mathcal{A} \times \mathcal{A}$, which we call the *information order*. Let $\pi_1, \pi_2 \in \text{Pre}(\mathcal{A})$. We write $\pi_1 \preceq \pi_2$ if π_1 is contained in π_2 as subsets of $\mathcal{A} \times \mathcal{A}$. The minimum element in the information order is the preference relation, written ϵ , defined $a \succeq a$ for all $a \in \mathcal{A}$. An agent with this preference relation has not compared any alternatives. The maximum element, written ι , in the information order satisfies $a \succeq b$ for all pairs $(a, b) \in \mathcal{A} \times \mathcal{A}$. An agent with this preference relation is indifferent between any two alternatives. The information order extends to preference profiles via the product. Suppose $\pi, \pi' \in \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$. Then, in abuse of notation, we write $\pi \preceq \pi'$ if $\pi_i \preceq \pi'_i$ for every $i \in \mathcal{N}$.

Preference relations satisfying Assumption 1 and 2 can be obtained from arbitrary binary relations via a unary operation. Suppose $\pi_1, \pi_2 \subseteq \mathcal{A} \times \mathcal{A}$ are arbitrary binary relations. Then, $\pi_2 \circ \pi_1 = \{(a, b) \in \mathcal{A} \times \mathcal{A} : \exists c, (a, c) \in \pi_1, (c, b) \in \pi_2\}$ is a preference relation that is the composition of π_1 and π_2 .

Definition 2 (Transitive Closure): Suppose $\pi \subseteq \mathcal{A} \times \mathcal{A}$. The *transitive closure* of π is the preference relation $\pi^* =$

$\bigcup_{p=1}^{\infty} \pi^{\circ p}$, where $\pi^{\circ k} = \overbrace{\pi \circ \pi \circ \dots \circ \pi}^k$. The *transitive-reflexive closure* of π is the preference relation $\pi^+ = (\pi \cup \epsilon)^*$.

It follows that if π is an arbitrary binary relation on \mathcal{A} , then π^+ is the smallest (transitive-reflexive) preference relation on \mathcal{A} containing π ; see, for instance, [27, Theorem 1.17].

D. Meets & Joins of Preferences

The transitive closure as well as the union and intersection of preference relations are enough to specify the structure of a lattice.

Theorem 1: $(\text{Pre}(\mathcal{A}), \succeq)$ is a complete lattice with meets and joins given by the following

$$\bigwedge_{j \in J} \pi_j = \bigcap_{j \in J} \pi_j, \quad \bigvee_{j \in J} \pi_j = \left(\bigcup_{j \in J} \pi_j \right)^+. \quad (2)$$

Proof: See Appendix. ■

We now interpret the meet and join operations of the information order. A preference $a \succeq b$ is in the meet $\pi_1 \wedge \pi_2$ if and only if it is in both π_1 and π_2 . Thus, the meet of two preference relations constitutes a consensus between them. The join is somewhat more subtle. For one, $\pi_1 \vee \pi_2$ is the smallest transitive-reflexive relation containing $\pi_1 \cup \pi_2$. In case that \mathcal{A} is finite, another interpretation exacts what it means for a pair-wise comparison to be in the join of two preference relations.

Proposition 1: Suppose \mathcal{A} is finite, and suppose $\pi_1, \pi_2 \in \text{Pre}(\mathcal{A})$. Then, $(a \succeq b) \in \pi_1 \vee \pi_2$ if and only if either, (i) $(a \succeq b) \in \pi_1 \cup \pi_2$, or, (ii) there exist a chain

$$a = c_0 \succ_{j_1} c_1 \succ_{j_2} \dots \succ_{j_{\ell-1}} c_{j_{\ell-1}} \succ_{j_\ell} c_\ell = b \quad (3)$$

such that $(c_{m-1} \succ_{j_m} c_m) \in \pi_1 \Delta \pi_2$, the symmetric difference of π_1 and π_2 .

Proof: Follows from Definition 2 and Theorem 1. ■

IV. PREFERENCE DYNAMICS

In this section, we introduce a preference dynamics model in a general message-passing framework and discuss particular cases that exploit the meet and join operation of the information order. We demonstrate a specific case of our model in Section VI.

A. A Message-Passing Model

We model how agents update their preferences during every round of interactions via a message-passing algorithm. The algorithm follows a familiar gather-scatter paradigm, commonly seen in graph machine learning [29], [30]. Agents send messages (also preference relations) to their neighbors in each round. Messages are generated by applying a function $\psi_i(\pi_i, \hat{\pi})$ to the preference relation π_i of the sending agent and an estimate $\hat{\pi}$ of the receiving agent’s preference relation. Agents collect the messages (preference relations) that they have received and aggregate them into a single preference relation with a function $\text{Aggregate}_i(\cdot)$. Finally, the agents update their preference (Line 9) as a result of the aggregation and their prior-held preference via a function $\varphi_i(\pi_i, \cdot)$. If agents update preferences with full synchrony, the algorithm is succinctly summarized by a map F_i in (1) of the form

$$F_i(\pi) = \varphi_i\left(\pi_i, \text{Aggregate}_i[\psi_j(\pi_j, \pi_i)]_{j \in \mathcal{N}_i}\right). \quad (4)$$

B. Discussion

For the remainder of this section, we discuss candidates for the functions ψ , $\text{Aggregate}(\cdot)$, and φ , noting how various personalities of agents affect the choice of functions. In particular, we discuss functions that utilize the lattice structure of the information order.

a) Messages: Message-passing algorithms (at least in the context of graph neural networks [31]) have been compared to diffusion processes. Message functions $\psi_i(\pi_i, \pi_j)$ that are invariant to the second argument are said to be *isotropic*, otherwise *anisotropic*. In the isotropic case, if $\psi_i(\pi_i) = \pi_i$, Agent i represents their preference faithfully and communicates it to Agent j . Such agents can be thought of as being honest. Dishonest isotropic agents may misrepresent their true preferences with a non-trivial map $\psi_i : \text{Pre}(\mathcal{A}) \rightarrow \text{Pre}(\mathcal{A})$. In the opinion dynamics literature [22], behaviors such as exaggerating opinions, restricting the topics of discussion, or lying can be represented by similar maps, e.g. a map ψ_i which reverses the relative ordering of every pair of alternatives. In the anisotropic case, agents represent their preferences as a result of feedback. For instance, Agent i could mimic the preference of Agent j , i.e. $\psi_i(\pi_i, \pi_j) = \pi_j$, but more complex feedback mechanisms are possible.

b) Aggregation: Aggregating preference relations accurately is a significant challenge. Fortunately, the lattice structure of the information order provides some reasonable choices. In our message-passing framework, we sub-index the neighbor set, i.e., $\mathcal{N}_i = \{j_1, j_2, \dots, j_n\}$. We first consider some favorable properties of aggregation functions. Suppose $\text{Aggregate}(\cdot) : \bigcup_{n=1}^N \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A}) \rightarrow \text{Pre}(\mathcal{A})$ is an aggregation function sending an n -tuple of preference relations to an aggregate preference relation, $\text{Aggregate}(\pi_{j_1}, \pi_{j_2}, \dots, \pi_{j_n})$. We say $\text{Aggregate}(\cdot)$ satisfies *anonymity* if for all n -tuples, $\text{Aggregate}(\pi_{j_1}, \dots, \pi_{j_n}) = \text{Aggregate}(\pi_{\sigma(j_1)}, \dots, \pi_{\sigma(j_n)})$ for all permutations $\sigma \in \Sigma_n$, *unanimity* if $\pi_j = \pi_{j'} = \pi$ for all $j, j' \in \{j_1, j_2, \dots, j_n\}$ implies $\text{Aggregate}(\pi_{j_1}, \pi_{j_2}, \dots, \pi_{j_n}) = \pi$, and *r -middle* if, given $1 \leq r \leq n$, the ordering $\pi_{j_1} \preceq \dots \preceq \pi_{j_r} \preceq \dots \preceq \pi_{j_n}$ implies $\text{Aggregate}(\pi_{j_1}, \dots, \pi_{j_n}) = \pi_{j_r}$.

We now propose a family of aggregation functions that satisfies all three properties. We suppose agents have different aggregation functions in this family, depending on the characteristics of the agent.

Definition 3: Suppose $r \geq 0$. The r -median is the aggregation function

$$\text{Median}_r(\pi_{j_1}, \pi_{j_2}, \dots, \pi_{j_n}) = \bigvee_{\substack{J \subseteq \{1, 2, \dots, n\} \\ |J| \geq r}} \left(\bigwedge_{m \in J} \pi_{j_m} \right).$$

Proposition 2: $\text{Median}_r(\cdot)$ satisfies the axioms of anonymity, unanimity, and r -middle.

Proof: For the first, the set $\{I \subseteq \{1, 2, \dots, n\} : |I| \geq r\}$ is invariant under permutation. For the second, use the idempotence property (Lemma 1). For the third, by anonymity, we can, without loss of generality, write each J in the set $\{I \subseteq \{1, 2, \dots, n\} : |I| \geq r\}$ as $\{j_m, j_{m+1}, \dots, j_{m+r+1}\}$. ■

Anonymity implies the order messages are received or the identity of the sender has no bearing on the aggregate preference, a notion of privacy and fairness, respectively. Unanimity implies that if every agent has the same preference, the aggregate should, surely, reflect this. Finally, the r -middle condition implies, if a series of agents contain each others preferences, meaning agents' preferences differ only by adding comparisons to the preference relation of an existing agent, the aggregate preference will be the r -th preference. In the following observation, the r -median can be interpreted as a coalition-formation mechanism.

Proposition 3: Suppose \mathcal{A} is finite. Then, $a \succsim b$ is in $\text{Median}_r(\pi_{\mathcal{N}_i})$ if and only if there exist a chain $a = c_0 \succsim_{j_1} c_1 \succsim_{j_2} \dots \succsim_{j_{\ell-1}} c_{j_{\ell-1}} \succsim_{j_\ell} c_\ell = b$ and a federation of neighbors into coalitions $\{J_m : |J_m| \geq r\}_{m=1}^\ell \subseteq \mathcal{N}_i$ such that $(c_{m-1} \succsim_{j_m} c_m) \in \pi_j$ for all $j \in J_m$.

Proof: Follows from Proposition 1. ■

Thus, among agents in the neighbor set \mathcal{N}_i , coalitions of agents, which agree on a particular comparison, incorporate that comparison into the aggregate preference relation. The value of r is the threshold for constituting a ‘‘majority rule,’’ the minimum number of agents needed to form a coalition. If r is small, only a few agents need to come to a consensus on a particular comparison. Similarly, an agent with large r , requires a larger coalition of agents to reach consensus. Agents, thus, have different values of r , depending on their level of stubbornness. In the extremal cases, if $r = n$, then the median is the join projection $\bigvee_{m=1}^n \pi_{j_m}$. On the other hand, if $r = 1$, the median is the meet projection $\bigwedge_{m=1}^n \pi_{j_m}$.

c) Updates: We examine four update functions which, along with the a choice of ψ and $\text{Aggregate}(\cdot)$, reflect the personality of an agent. Suppose $\pi_{\text{agg}} = \text{Aggregate}_i[\psi_j(\pi_j, \pi_i)]_{j \in \mathcal{N}_i}$ is the result of aggregating messages. Then, an agent that is single-minded will not take into account the preferences of neighbors, i.e., $\varphi_i(\pi_i, \pi_{\text{agg}}) = \pi_i$, the *prior update* rule. On the other hand, an agent who is unusually open-minded might ignore their prior preferences, i.e., $\varphi_i(\pi_i, \pi_{\text{agg}}) = \pi_{\text{agg}}$, the *posterior update* rule. The choice of meet or join for φ has interesting implications as well. Suppose an agent is examining the validity of a prior $a \succsim b$. If that comparison is in π , then the *meet update*, i.e., $\varphi_i(\pi_i, \pi_{\text{agg}}) = \pi_i \wedge \pi_{\text{agg}}$, reflects a confirmation bias. On the other hand, if an agent is to learn new comparisons by observing other agents' preferences, the join, i.e., $\varphi_i(\pi_i, \pi_{\text{agg}}) = \pi_i \vee \pi_{\text{agg}}$, models the process of integrating a prior-held preference with new observations.

V. STABLE PREFERENCE PROFILES

Stable preference profiles are those fixed in the preference dynamical system, e.g., the message-passing model in Section IV. In this section, we study the structure of stable preferences as well as provide sufficient conditions for when an initial preference profile will converge to a stable preference profile.

Suppose $F = (F_1, F_2, \dots, F_N)$ is given, i.e., $F : \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A}) \rightarrow \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$ is an arbitrary iterative map sending preference profiles to preference profiles. The *equilibrium points* of the resulting global preference dynamical

system

$$\begin{aligned}\pi(t+1) &= F(\pi(t)) \\ \pi(0) &= \pi_0\end{aligned}\quad (5)$$

are collected in the following set $\mathcal{S} = \{\pi \in \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A}) : F(\pi) = \pi\}$, precisely the fixed points of F .

A. Existence & Structure of Equilibrium Points

We constrain the function class of iterative maps F in (5) to the class of monotone maps which, we argue, lend themselves to modeling preference updates. Monotone maps preserve the product information order on preference profiles. A counterfactual interpretation is in order. If an agent were to reveal more preferences not contradicting existing ones, then the preference relation held in the next round would only contain more comparisons than if the agent decided to not reveal the additional preferences. The following lemma discusses conditions under which the iterative maps F defined in Section IV are monotone.

Lemma 2: Suppose, for each $i \in \mathcal{N}$, the message-passing map ψ_i is isotropic and monotone, and suppose the aggregation function $\text{Aggregate}_i(\cdot)$ as well as the update function φ_i are (evaluations of) lattice polynomials (see Appendix). Then, the iterative map $F = (F_1, F_2, \dots, F_N)$ defined by the composition of these maps (4) is monotone.

Proof: Application of Lemma 4 (see Appendix) and the fact that the composition and product of monotone functions is monotone. ■

The usual approach to guarantee the existence of equilibrium points in discrete-time systems is to guarantee the existence of fixed points of the iterative map. However, standard approaches to proving the existence of fixed points, e.g., Brouwer’s fixed-point theorem, require, at a minimum, continuity of the iterative map, which we do not have here. The following fixed point theorem, due to Tarski [32] and Knaster [33], guarantees the existence of fixed points without continuity.

Lemma 3 (Tarski Fixed Point Theorem [34]): Suppose \mathcal{L} is a complete lattice and $f : \mathcal{L} \rightarrow \mathcal{L}$ is monotone. Then, the fixed point set $\text{Fix}(f)$ is a complete lattice.

Proof: See [27, Theorem 12.2]. ■

The Tarski Fixed Point Theorem does more than guarantee the existence of fixed points, including a minimum and maximum fixed point which coincide precisely when there is exactly one fixed point. It also sheds light on the structure of the fixed points. We apply the Tarski Fixed Point Theorem to characterize the equilibrium points of the preference dynamics (5) with a monotone iterative map.

Theorem 2: Suppose $\pi(t)$ evolves according to (5) with iterative map F , and suppose F is monotone. Then, the solutions \mathcal{S} form a complete lattice.

Proof: By Theorem 1, $\text{Pre}(\mathcal{A})$ is complete. It follows that $\prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$ is complete: the finite product of complete lattices is complete. Apply Lemma 3. ■

In particular, Theorem 2, together with Lemma 2, implies that the message-passing preference mechanisms discussed in Section IV have stable preference profiles. Thus, equilibrium

points of preference dynamics exist, although are not, in general, unique. However, by virtue of the fixed points forming a lattice, there exist unique greatest and least equilibrium points with respect to the product information order. The greatest (resp. least) equilibrium point is the stable preference profile with the most (resp. least) alternatives compared. More generally, the meet and join of stable preference profiles exist: if π, π' are stable preference profiles, there is a largest *stable* preference profile $\pi \wedge_{\mathcal{S}} \pi'$ contained by both π and π' . There also exists a smallest *stable* preference profile $\pi \vee_{\mathcal{S}} \pi'$ containing both π and π' .

B. Convergence to Equilibrium Points

It is well-known that the standard proof of the Tarski Fixed Point Theorem is “non-constructive” [35] in that it does not explicitly produce a fixed point, only guarantees the existence of a fixed point. Thus, Theorem 2 is only a result about the existence and structure of stable preferences, not a method to compute them. In the remainder of this section, we consider iterative maps F which guarantee the convergence of $\pi(0)$ to some stable preference profile $\pi \in \mathcal{S}$. We say $\pi(t)$ *converges in finite time* to π if there exists $t_0 \geq 0$ such that $\pi(t) = \pi$ for all $t \geq t_0$. In the following two results, we make the assumption that \mathcal{A} is finite, although it is possible (though technical) to also reason about convergence in the case that \mathcal{A} is not finite (for instance, see [27, Theorem 12.9]).

Theorem 3: Suppose $\pi(t)$ evolves according to (5) with iterative map F , suppose \mathcal{A} is finite, and suppose F is inflationary. Given an initial preference profile $\pi(0) \in \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$, let $\pi^* = \bigvee_{t \geq 0} \pi(t)$. Then, $\pi^* \in \mathcal{S}$, and $\pi(t)$ converges to π^* in finite time.

Proof: Let $\pi(0) \in \prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$, and suppose $F(\pi) \succcurlyeq \pi$ for all π . In particular, $F(\pi(0)) \succcurlyeq \pi(0)$, $F(F(\pi(0))) \succcurlyeq F(\pi(0))$, and so on. Thus, trajectories form an ascending chain $\dots \succcurlyeq F^t(\pi(0)) \succcurlyeq \dots \succcurlyeq F(\pi(0)) \succcurlyeq \pi(0)$. Because $\prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$ is a complete lattice, the quantity $\pi^* = \bigvee_{t \geq 0} F^t(\pi(0))$ exists. Because \mathcal{A} is finite, and, thus, the lattice $\prod_{i \in \mathcal{N}} \text{Pre}(\mathcal{A})$ is finite, $\pi^* \in \{F^t(\pi(0))\}_{t=0}^{\infty}$, i.e., the chain eventually terminates at π^* . Therefore, there exist $t_0 \geq 0$ such that $F^{t_0}(\pi(0)) = F^{t_0+t}(\pi(0)) = \pi^*$ for all $t \geq t_0$ in the ascending chain above. ■

A similar argument and result holds for F deflationary. In particular, the join (resp. meet) update, $\varphi(\pi, \pi_{\text{agg}}) = \pi \vee \pi_{\text{agg}}$ (resp. $\varphi(\pi, \pi_{\text{agg}}) = \pi \wedge \pi_{\text{agg}}$) satisfies the inflationary (resp. deflationary) condition locally. In general, if F is inflationary, then, for all $i \in \mathcal{N}$, $\pi_i(t+1)$ contains every comparison in $\pi_i(t)$. Thus, comparisons are added to each agent’s preferences each round, but never removed. On the other hand, if F is deflationary, comparisons are removed from each agent’s preferences each round, but never added. In future work, we want to relax the assumptions of Theorem 3 to encompass preference dynamics in which some agents may add comparisons each round, while others may remove them.

VI. EXPERIMENTS

In this section, we perform a simulation of the preference dynamics model proposed in Section IV and validate our theoretical findings in Section V. We measure the convergence

of preference profiles to equilibria using the Kendall tau [36] (also known as the Kemeny-Snell-Bogart [37]) metric, defined $d(\pi_1, \pi_2) = \#\{(a, b) : (a \succ b) \in \pi_1, (b \succ a) \in \pi_2\}$ for preference relations $\pi_1, \pi_2 \in \text{Pre}(\mathcal{A})$. This distance is a measure of disagreement between two agents, where agents are said to disagree if they hold opposite comparisons. Given a preference profile π and a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, consider the quantity $\mathcal{D}(\pi) = \sum_{(i,j) \in \mathcal{E}} d(\pi_i, \pi_j)$. Because the dynamical system (5) resembles a diffusion process, we instill the term *Dirichlet energy* to refer to the quantity $\mathcal{D}(\cdot)$. The Dirichlet energy quantifies the total disagreement between agents who are connected in \mathcal{G} .

A. Experimental Setup

Fixing the number of agents and neighbors of each agent, we assign message ψ , aggregation $\text{Aggregate}(\cdot)$, and update φ functions to every agent; we set φ and ψ to be the same rule for every agent, but let $\text{Aggregate}(\cdot)$ vary from agent to agent. We assume agents are truthful, i.e., $\psi(\pi_i, \pi_j) = \pi_i$, and we assume agents update via the join operation, i.e., $\varphi(\pi_i, \pi_{\text{agg}}) = \pi_i \vee \pi_{\text{agg}}$. We let $\text{Aggregate}(\cdot) = \text{Median}_r(\cdot)$, and, modeling agents with different levels of stubbornness, assign different values of $r \in \{1, 2, \dots, k\}$ to agents randomly. Thus, the iterative map in the global dynamics (5) is monotone and inflationary. Next, we choose several initial preference profiles over the alternative set $\mathcal{A} = \{1, 2, 3, 4, 5\}$. Each initial profile $\pi(0)$ is given by a choice of preference relation for each agent, which, in turn, is produced by sampling possible edges, selecting each $a \succ b$ with probability p , and discarding preference relations that violate transitivity. Then, we randomly generate k -regular graphs with nodes $\mathcal{N} = \{1, 2, \dots, N\}$ for several values of k , using an algorithm introduced by [38]. We compute the trajectories of the preference dynamics equation (5) for the initial preference profiles we selected and, at every round $t = 0, \dots, t_{\text{max}}$, we also compute the Kendall tau distance $d(\pi_i(t), \pi_j(t))$ for every $(i, j) \in \mathcal{E}$ and the Dirichlet energy $\mathcal{D}(\pi(t))$.

B. Discussion of Results

We considered $N = 20$ agents and compared the convergence of their preference profiles for various k -regular graphs. As expected, every trajectory converged to a stable preference profile (see Theorem 3). We make some additional observations. The Dirichlet energy, for each trajectory, was non-decreasing, reflecting a greater or equal level of disagreement each round. We were surprised that trajectories tended to cluster into stable preference profiles with a similar level of total disagreement. One possible explanation for this could be that several profiles in the lattice of stable profiles (see Theorem 1) have similar levels of Dirichlet energy. We examine the emergent behavior of disagreement more closely by comparing the Kendall tau distances between connected agents in the initial profile $\pi(0)$ versus the final stable profile $\pi(t_{\text{max}})$. We display our results with a heat map over the network (see Fig. 1). We observe, at least in this example, that agents, who initially have some disagreement, end up disagreeing more, while agents who don't disagree

initially, never disagree. Hence, it appears, the dynamics make disagreements more pronounced.

VII. CONCLUSION

We introduced a mathematical framework of lattice dynamical systems on networks and used it to design and analyze a preference dynamics model in which agents reject or accept comparisons between alternatives in a majoritarian fashion. The mechanism we introduced can lead to either dissensus or consensus in the steady-state, depending on the distribution of the r -values of agents which reflect their level of stubbornness. In the future, we want to incorporate control inputs, e.g. agents whose preferences can be controlled, into our model.

REFERENCES

- [1] S. L. Lingel, M. Sargent, T. R. Gulden, T. McDonald, and P. Rockstroh, *Leveraging Complexity in Great-Power Competition and Warfare*. RAND Corporation, 2021.
- [2] M. C. Munger, *Choosing in groups: Analytical politics revisited*. Cambridge University Press, 2015.
- [3] S. Natarajan and P. Tadepalli, "Dynamic preferences in multi-criteria reinforcement learning," in *Proceedings of the 22nd international conference on Machine learning*, pp. 601–608, 2005.
- [4] T. Roughgarden, "Algorithmic game theory," *Communications of the ACM*, vol. 53, no. 7, pp. 78–86, 2010.
- [5] D. M. Ziegler, N. Stiennon, J. Wu, T. B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving, "Fine-tuning language models from human preferences," *arXiv preprint arXiv:1909.08593*, 2019.
- [6] A. F. Beardon, J. C. Candeal, G. Herden, E. Induráin, and G. B. Mehta, "The non-existence of a utility function and the structure of non-representable preference relations," *Journal of Mathematical Economics*, vol. 37, no. 1, pp. 17–38, 2002.
- [7] K. Eliaz and E. Ok, "Indifference or indecisiveness? choice-theoretic foundations of incomplete preferences," *Games and Economic Behavior*, vol. 56, pp. 61–86, 2006.
- [8] J. C. Harsanyi, "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of political economy*, vol. 63, no. 4, pp. 309–321, 1955.
- [9] S. O. Hansson, "Changes in preferences," *Theory and Decision*, vol. 38, no. 1, pp. 1–28, 1995.
- [10] G. J. Stigler and G. S. Becker, "De gustibus non est disputandum," *The american economic review*, vol. 67, no. 2, pp. 76–90, 1977.
- [11] H. Riess and R. Ghrist, "Diffusion of information on networked lattices by gossip," in *2022 IEEE Conference on Decision and Control (CDC)*, (Cancun, Mexico), 2022.
- [12] H. Riess, M. Munger, and M. M. Zavlanos, "Max-plus synchronization in decentralized trading systems," *arXiv preprint arXiv:2304.00210*, 2023.
- [13] G. Currello and L. Sinander, "The preference lattice," *arXiv preprint arXiv:1902.07260*, 2019.
- [14] F. Karacal and R. Mesiar, "Aggregation functions on bounded lattices," *International Journal of General Systems*, vol. 46, no. 1, pp. 37–51, 2017.
- [15] C. P. Chambers and A. D. Miller, "Rules for aggregating information," *Social Choice and Welfare*, vol. 36, no. 1, pp. 75–82, 2011.
- [16] K. J. Arrow, *Social choice and individual values*, vol. 12. Yale university press, 2012.
- [17] J.-P. Barthélemy and M. F. Janowitz, "A formal theory of consensus," *SIAM Journal on Discrete Mathematics*, vol. 4, no. 3, pp. 305–322, 1991.
- [18] B. Jean-Pierre, L. Bruno, and M. Bernard, "On the use of ordered sets in problems of comparison and consensus of classifications," *Journal of Classification*, 1986.
- [19] H. Attiya, M. Herlihy, and O. Rachman, "Atomic snapshots using lattice agreement," *Distributed Computing*, vol. 8, pp. 121–132, 1995.
- [20] H. Noorazar, "Recent advances in opinion propagation dynamics: A 2020 survey," *The European Physical Journal Plus*, vol. 135, pp. 1–20, 2020.
- [21] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical association*, vol. 69, no. 345, pp. 118–121, 1974.




Fig. 1: On a k -regular graph ($k = 4$) with $N = 20$ nodes with a randomly-selected initial preference profile, we plot (a) the Kendall tau distance $d(\pi_i(t), \pi_j(t))$ for every $(i, j) \in \mathcal{E}$ for both $t = 0$ (initial disagreement) and $t = 15$ (final disagreement) with a heat-map on the edges of the graph. We also plot (b) the Dirichlet energy of the trajectory (black) as well as the energies of other trajectories from different initial conditions (blues).

- [22] J. Hansen and R. Ghrist, “Opinion dynamics on discourse sheaves,” *SIAM Journal on Applied Mathematics*, vol. 81, no. 5, pp. 2033–2060, 2021.
- [23] J. Ghaderi and R. Srikant, “Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate,” *Automatica*, vol. 50, no. 12, pp. 3209–3215, 2014.
- [24] R. Hegselmann and U. Krause, “Opinion dynamics and bounded confidence intervals,” *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.
- [25] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, “On krause’s multi-agent consensus model with state-dependent connectivity,” *IEEE Transactions on Automatic Control*, 2009.
- [26] M. Hayhoe, F. Alajaji, and B. Gharesifard, “A polya contagion model for networks,” *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 1998–2010, 2017.
- [27] S. Roman, *Lattices and ordered sets*. Springer, 2008.
- [28] G. Birkhoff, *Lattice theory*, vol. 25. American Mathematical Soc., 1940.
- [29] A. Dudzik and P. Veličković, “Graph neural networks are dynamic programmers,” *arXiv preprint arXiv:2203.15544*, 2022.
- [30] A. Dudzik, T. von Glehn, R. Pascanu, and P. Veličković, “Asynchronous algorithmic alignment with cocycles,” *arXiv preprint arXiv:2306.15632*, 2023.
- [31] S. A. Tailor, F. Opolka, P. Lio, and N. D. Lane, “Do we need anisotropic graph neural networks?,” in *International Conference on Learning Representations*, 2021.
- [32] A. Tarski, “On the calculus of relations,” *The Journal of Symbolic Logic*, vol. 6, no. 3, pp. 73–89, 1941.
- [33] B. Knaster, “Un theoreme sur les fonctions d’ensembles,” *Ann. Soc. Polon. Math.*, vol. 6, pp. 133–134, 1928.
- [34] A. Tarski, “A lattice-theoretical fixpoint theorem and its applications,” *Pacific journal of Mathematics*, vol. 5, no. 2, pp. 285–309, 1955.
- [35] P. Cousot and R. Cousot, “Constructive versions of tarski’s fixed point theorems,” *Pacific journal of Mathematics*, vol. 82, no. 1, pp. 43–57, 1979.
- [36] M. G. Kendall, “A new measure of rank correlation,” *Biometrika*, vol. 30, no. 1/2, pp. 81–93, 1938.
- [37] H. Nishimura and E. A. Ok, “A class of dissimilarity semimetrics for preference relations,” *Mathematics of Operations Research*, 2023.
- [38] J. H. Kim and V. H. Vu, “Generating random regular graphs,” in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pp. 213–222, 2003.

APPENDIX

A. Lattice Polynomials

Given variables $\pi_1, \pi_2, \dots, \pi_n$, a *lattice polynomial* p is a term in the language $\pi_1, \pi_2, \dots, \pi_n \mid \pi \wedge \pi' \mid \pi \vee \pi'$, i.e., an expression formed by finite application of these symbols and parentheses. Suppose \mathcal{L} is lattice. Then, a lattice polynomial p defines an evaluation map $\text{ev}_p : \mathcal{L}^n \rightarrow \mathcal{L}$ by substituting an element of \mathcal{L} into each variable of p .

Lemma 4: Suppose p is a lattice polynomial. Then, $\text{ev}_p(\pi_1, \pi_2, \dots, \pi_n)$ is monotone in the product lattice \mathcal{L}^n .

Proof: See [28, §II.5 Lemma 4]. ■

B. Closure Operators & Systems

Suppose \mathcal{X} is a set. A *closure operator* is a monotone, inflationary map $\text{cl} : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$ satisfying the additional property: $\text{cl}(\text{cl}(S)) = \text{cl}(S)$ for all $S \subseteq \mathcal{X}$. A *closure system* is a collection $\mathcal{F} = \{S_\alpha\}_{\alpha \in I}$ of subsets of \mathcal{X} such that $\bigcap_{\alpha \in I'} S_\alpha \in \mathcal{F}$ for every subcollection $\mathcal{F}' = \{S_\alpha\}_{\alpha \in I' \subseteq I}$. It follows that a closure system is a complete lattice, and there is a correspondence between closure operators and closure systems given by

$$\begin{aligned} \text{cl} &\mapsto \text{Cl}(\mathcal{X}) \triangleq \{S \subseteq \mathcal{X} : \text{cl}(S) = S\}, \\ \mathcal{F} &\mapsto \left(\text{cl} : S \mapsto \bigcap \{S' \in \mathcal{F} : S \subseteq S'\} \right). \end{aligned} \quad (6)$$

Lemma 5: The correspondence in (6) are inverse bijections. Furthermore, $\text{Cl}(\mathcal{X})$ is a complete lattice with the join and meet operations for an arbitrary collection of subsets $\mathcal{U} \subseteq \text{Cl}(\mathcal{X})$,

$$\begin{aligned} \bigvee_{\text{Cl}(\mathcal{X})} \mathcal{U} &= \text{cl} \left(\bigcup \{S : S \in \mathcal{U}\} \right), \\ \bigwedge_{\text{Cl}(\mathcal{X})} \mathcal{U} &= \bigcap \{S : S \in \mathcal{U}\}. \end{aligned} \quad (7)$$

Proof: See [27, Theorem 3.7-3.8]. ■

C. Proof of Theorem 1

We first show that the transitive-reflexive closure is a closure operator on the set $\mathcal{X} = \mathcal{A} \times \mathcal{A}$.

Lemma 6: Suppose \mathcal{A} is an arbitrary set. The transitive-reflexive closure (Definition 2) of binary relations on \mathcal{A} is a closure operator.

Proof: Suppose $\pi \subseteq \mathcal{A} \times \mathcal{A}$. Then, it is straightforward to check $\pi^{++} = \pi^+$, i.e., the transitive-reflexive closure of a transitive-reflexive relation is transitive-reflexive, and $\pi^+ \supseteq \pi$, i.e., the transitive-reflexive closure of a relation contains the relation. For monotonicity, it suffices to check that $\pi_1 \subseteq \pi_2$ implies $\pi_1 \circ \pi_1 \subseteq \pi_2 \circ \pi_2$. ■

Applying Lemma 5 and Lemma 6 yields our result.