

The Semi-random Likelihood of Doctrinal Paradoxes

Ao Liu and Lirong Xia

Department of Computer Science, Rensselaer Polytechnic Institute
110 8th St, Troy, NY 12180, USA
liua6@rpi.edu, xial@cs.rpi.edu

Abstract

When aggregating logically interconnected judgements from n agents, the result might be logically inconsistent. This phenomenon is known as the doctrinal paradox, which plays a central role in the field of judgement aggregation. Previous work has mostly focused on the worst-case analysis of the doctrinal paradox, leading to many impossibility results. Little is known about its likelihood of occurrence in practical settings, except for the study under certain distributions by List in 2005. In this paper, we characterize the likelihood of the doctrinal paradox under a general and realistic model called semi-random social choice framework (proposed by Xia in 2020). In the framework, agents' ground truth judgements can be arbitrarily correlated, while the noises are independent. Our main theorem states that under mild conditions, the semi-random likelihood of the doctrinal paradox is either 0, $\exp(-\Theta(n))$, $\Theta(n^{-0.5})$ or $\Theta(1)$. This not only answers open questions by List in 2005, but also draws clear lines between situations with frequent paradoxes and with vanishing paradoxes.

Introduction

The field of *judgement aggregation* studies the aggregation of agents' judgements (yes/no votes) on multiple logically interconnected propositions in a consistent way. For example, suppose a defendant is involved in a traffic accident where a victim is injured. A jury of n agents (jurors) are asked to make collective decisions on the following three propositions:

Proposition ω_1 : The injury is caused by the defendant.

Proposition ω_2 : The injury is serious enough.

Proposition ω_3 : The defendant is guilty.

Each agent is asked to provide a yes/no judgement to each proposition in a logically consistent way, i.e., her answer to ω_3 is Y if and only if her answers to both ω_1 and ω_2 are Y, or equivalently, $\omega_3 \leftrightarrow \omega_1 \wedge \omega_2$ must hold. Then, agents' judgements, called a *profile*, are aggregated by a rule r .

One classical rule for judgement aggregation is *proposition-wise majority*, which performs majority voting on each proposition separately. Unfortunately, the output of proposition-wise majority can be logically

inconsistent. This phenomenon is known as the *doctrinal paradox*, as shown in the following example.

Example 1 (Doctrinal paradox). *Suppose there are three agents ($n = 3$), whose judgements on ω_1 , ω_2 , and ω_3 are shown in Table 1. Notice that while every agent's judgements are consistent with the logic relationship $\omega_3 \leftrightarrow \omega_1 \wedge \omega_2$, the output of proposition-wise majority is inconsistent with the logic.*

Propositions	ω_1	ω_2	ω_3	$\omega_3 \stackrel{?}{\leftrightarrow} \omega_1 \wedge \omega_2$
Agent 1	Y	N	N	True
Agent 2	N	Y	N	True
Agent 3	Y	Y	Y	True
Aggregation	Y	Y	N	False

Table 1: Doctrinal paradox.

In general, the doctrinal paradox can happen for $p \geq 1$ *premises* (logically independent propositions), plus one *conclusion*, whose value is determined by the values of the premises. The doctrinal paradox “*originated the whole field of judgement aggregation*” (Grossi and Pigozzi 2014), and continued to play a central role since then (List 2012; Grossi and Pigozzi 2014; Brandt et al. 2016). A large body of literature has shown the negative effects of the doctrinal paradox in law (Kornhauser and Sager 1986; Hanna 2009; Chilton and Tingley 2012), economics (Mongin 2019), computational social choice (Bonneton 2010; Brandt et al. 2016), philosophy (Sorensen 2003; Mongin 2012) and psychology (Bonneton 2011), etc. Unfortunately, the paradox is unavoidable under mild assumptions (List and Pettit 2002), and there is a large body of literature on impossibility theorems as well as ways for escaping from them (List and Polak 2010; Grossi and Pigozzi 2014; Endriss 2016).

Despite the mathematical impossibilities, it is unclear how frequently doctrinal paradoxes happens in reality. Empirical studies have been limited by the lack of datasets, as many classical applications of judgement aggregations are in high-stakes, low frequent, privacy-sensitive domains, such as judicial systems. To embrace AI-powered e-democracy, we believe that it is now the right time to theoretically study the likelihood of the paradox under realistic models. A low likelihood is desirable, because it suggests that, despite the fact that the paradox exists in the worst case, it is likelihood not be a big concern in practice.

While this approach is highly popular in voting (Nurmi 1999; Gehrlein and Lepelley 2011; Diss and Merlin 2021), it has not received due attention in judgement aggregation. The only known exception is the paper by List (2005), who provided sufficient conditions for the likelihood of doctrinal paradox to converge to 1 and to 0, respectively, under some distributions for two or three premises. Therefore, the following research question remains largely open:

How likely does the doctrinal paradox occur under realistic models?

Successfully addressing this question faces two challenges. The first challenge is *modeling*, i.e., it is unclear what probabilistic models are realistic. A large body of literature in social choice focused on the i.i.d. uniform distribution, known as *Impartial Culture (IC)*, which has been widely criticized of being unrealistic (Nehama 2011, 2013; Nurmi 1999; Brandl, Brandt, and Seedig 2016; Van Deemen 2014; Lehtinen and Kuorikoski 2007). The second challenge is the *technical hardness* in mathematically analyzing the likelihood. In fact, even characterizing the likelihood of doctrinal paradox under i.i.d. distributions is already highly challenging, and List’s 2005 work (List 2005) left many i.i.d. distributions as open questions, including IC.

Our Model. We adopt the *semi-random social choice* framework (Xia 2020, 2021b; Xia and Zheng 2021) to address the modeling challenge. The framework was inspired by the celebrated semi-random complexity analysis (Spielman and Teng 2009; Blum and Spencer 1995), and is more realistic and general than i.i.d. distributions, especially IC. We consider the setting with $p \geq 1$ premises and one conclusion determined by the logical connection function f , where every agent’s vote is a random variable whose distribution is in a set Π over all 2^p judgements on premises (See the Preliminaries Section for the formal definitions of distributions over votes). Then, given the aggregation rule r , a max-adversary aims at maximizing the likelihood of doctrinal paradox by setting agents’ distributions from Π . This leads to the definition of the *max-semi-random likelihood* of doctrinal paradox as follows.

$$\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n) \triangleq \sup_{\vec{\pi} \in \Pi^n} \Pr_{P \sim \vec{\pi}} (P \text{ is a doctrinal paradox}),$$

where $\vec{\pi} = (\pi_1, \dots, \pi_n) \in \Pi^n$ is the collection of the distributions of n agents chosen by the adversary, and profile P is the collection of n agents’ votes generated from $\vec{\pi}$. Similarly, the *min-semi-random likelihood* of doctrinal paradox is defined as follows.

$$\widetilde{\text{DP}}_{\Pi,r,f}^{\min}(n) \triangleq \inf_{\vec{\pi} \in \Pi^n} \Pr_{P \sim \vec{\pi}} (P \text{ is a doctrinal paradox}).$$

Notice that in both definitions, agents’ “ground truth” preferences, represented by $\vec{\pi}$, are arbitrarily correlated (the noises are still independent). The max-(respectively, min-) semi-random likelihood corresponds to the worst-case (best-case) analysis. A low max-semi-random likelihood is positive news, because it states that the paradox is unlikely to happen regardless of the adversary’s choice. A high min-semi-random likelihood is strongly negative news, because it states that the paradox is likely to happen with

non-negligible probability, regardless of agents’ ground truth preferences.

Our Contributions. In this paper, we prove the following general characterization of the semi-random likelihood of doctrinal paradox.

Theorem 1. (Semi-random Likelihood of Doctrinal Paradox, Informal.) *Under mild assumptions, for any $n \in \mathbb{N}$, any quota rule r , and any logical connection f , $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ is either 0, $\exp(-\Theta(n))$, $\Theta(n^{-1/2})$ or $\Theta(1)$, and $\widetilde{\text{DP}}_{\Pi,r,f}^{\min}(n)$ is either 0, $\exp(-\Theta(n))$, $\Theta(n^{-1/2})$ or $\Theta(1)$.*

The formal statement of the theorem also characterizes the condition for each case to happen. As commented above, the first three cases (0, $\exp(-\Theta(n))$, and $\Theta(n^{-1/2})$) of $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ are positive news, because they state that the doctrinal paradox vanishes as $n \rightarrow \infty$, see e.g., Example 7.

The last case ($\Theta(1)$) of $\widetilde{\text{DP}}_{\Pi,r,f}^{\min}(n)$ is negative news. Notably, a direct corollary of Theorem 1 (Corollary 2) characterizes the likelihood of doctrinal paradox under i.i.d. distributions that are not covered by List (2005), in particular the uniform distribution which naturally corresponds to IC in social choice. Experiments on synthetic data in the next Section confirm our theoretical results.

Proof techniques. The proof of Theorem 1 addresses technical hardness in (List 2005) by first converting the semi-random doctrinal paradox to a union of polyhedra, and then applying (Xia 2021a, Theorem 2). Such applications can be highly non-trivial as commented in (Xia 2021a), which we believe to be the case of this paper and are our main technical contribution.

Related Work and Discussions. Our work extends the results by List (2005) on i.i.d. distributions in three dimensions. First, we characterize the semi-random likelihood of doctrinal paradox, which is more general and realistic than i.i.d. distributions. Second, our theorem works for arbitrary number of premises and arbitrary logic connection function, while List (2005)’s results only work for two or three premises and conjunctive and disjunctive functions. Third, our theorem works for arbitrary quota rules, which extends the proposition-wise majority rule in (List 2005). As discussed above, as our theorem is a full characterization, a straightforward corollary of our theorem addresses the likelihood of doctrinal paradox under all i.i.d. cases that are left open in (List 2005).

The modern formulation of doctrinal paradox was due to Kornhauser and Sager (1986), though a similar paradox was described by Poisson in 1937 (Poisson 1837) and by Vacca in 1921 (Vacca 1921), see (Grossi and Pigozzi 2014). The paradox is closely related to the more general *discursive dilemma* (Pettit 2001), where agents also “vote” on the logic relationship as in the last column in Table 1. Doctrinal paradox is also a special case of responsibility voids, where no individual in the committee can morally be responsible voting outcome (Braham and VanHees 2011).

List and Pettit (2002) proved the first impossible theorem about doctrinal paradox, which has lead to many followup impossibility theorems (Pauly and Hees 2006; Mongin 2008; Dietrich and List 2008; Awad et al. 2017; Mongin 2019; Baharad, Neeman, and Rubinchik 2020; Marcoci and Nguyen 2020). Methods to escape from the paradox have been explored (Nehring and Pivato 2021; Rahwan and Tohmé 2010; Nehring and Puppe 2008; Lyon and Pacuit 2013). The connection between judgement aggregation and belief merging was established (Everaere, Konieczny, and Marquis 2017). Computational and machine learning aspects of judgement aggregation have also been explored, see, e.g., (Endriss, Grandi, and Porello 2012; Zhang and Conitzer 2019; Endriss et al. 2019; Baumeister et al. 2020).

Preliminaries

Basics of Judgement Aggregation. Let $[n] \triangleq \{1, \dots, n\}$ denote the set of n agents. Let p denote the number of premises. The binary judgement of agent j on the i -th premise is denoted by $\omega_{j,i} \in \{0, 1\}$, where $\omega_{j,i} = 1$ means ‘‘Yes’’ and $\omega_{j,i} = 0$ means ‘‘No’’. The conclusion is often denoted by ϕ , and sometimes by ω_{p+1} for simplicity. Because all premises are binary, there are $m = 2^p$ different combinations of *judgements* on the premises. To better present the results, instead of asking the agents to submit their judgements $\vec{\omega}_j = (\omega_{j,1}, \dots, \omega_{j,p})$, equivalently, we ask each agent to submit an m -dimensional vector \vec{v}_j , called her *vote*, whose $\vec{\omega}_j$ -th component is 1 and all other components are 0’s. Let $\mathcal{V} \triangleq \{\vec{v} \in \{0, 1\}^m : \|\vec{v}\|_1 = 1\}$ denote the set of all votes. For example, in the setting of Example 1, we have $p = 2$, $m = 4$, and the four judgements are $((0, 0), (0, 1), (1, 0), (1, 1))$. Agent 1’s judgements are $(1, 0)$, which means that her vote is $(0, 0, 1, 0)$, where the third component corresponds to the weight on judgements $(1, 0)$.

A probability distribution over the 2^p judgements is also called a *fractional vote*. Let $\mathcal{V}_{\text{frac}} \triangleq \{\vec{v} \in [0, 1]^m : \|\vec{v}\|_1 = 1\}$ denote the set of all fractional votes. Clearly, $\mathcal{V} \subseteq \mathcal{V}_{\text{frac}}$. For any fractional vote \vec{v} and any judgement $\vec{\omega}$, we let $\vec{v}(\vec{\omega})$ denote the weight (probability) on $\vec{\omega}$. For example, agent 2’s vote in Table 2 below is a fractional vote with 0.5 weight on each of $(1, 0)$ and $(1, 1)$.

Let $f : \{0, 1\}^p \rightarrow \{0, 1\}$ denote the *logical connection* between the premises and the conclusion. That is, for any judgements $\vec{\omega} \in \{0, 1\}^p$, $f(\vec{\omega})$ is the conclusion that is consistent with the logic. The domain of f can be naturally extended to (non-fractional) votes \mathcal{V} . That is,

$$\forall j \in [n], \phi_j = f(\vec{v}_j) = f(\vec{\omega}), \text{ where } \vec{v}_j(\vec{\omega}) = 1.$$

For example, in the setting of Example 1, the conclusion is 1 if and only if both premises are 1. Therefore,

$$f(\vec{v}) = \begin{cases} 1 & \text{if } \vec{v} = (0, 0, 0, 1) \\ 0 & \text{otherwise} \end{cases}$$

For any $1 \leq i \leq p + 1$, let Ω_i denote the set of all judgements whose i -th proposition is 1. That is,

$$\Omega_i \triangleq \begin{cases} \{\vec{\omega} \in \{0, 1\}^p : \vec{\omega}(i) = 1\} & \forall i \in [p] \\ \{\vec{\omega} \in \{0, 1\}^p : f(\vec{\omega}) = 1\} & \text{if } i = p + 1 \end{cases}$$

For example, in the setting of Example 1, $\Omega_1 = \{(1, 0), (1, 1)\}$, $\Omega_2 = \{(0, 1), (1, 1)\}$, and $\Omega_3 = \{(1, 1)\}$.

For any $n \in \mathbb{N}$, let $P = (\vec{v}_1, \dots, \vec{v}_n) \in (\mathcal{V}_{\text{frac}})^n$ denote a (*fractional*) *profile* of n agents. A (*judgement*) *aggregation rule* is a function $r : (\mathcal{V}_{\text{frac}})^n \rightarrow \{0, 1\}^{p+1}$, which takes a profile as input and outputs binary values for all premises and the conclusion. For any (fractional) profile $P \in (\mathcal{V}_{\text{frac}})^n$, we define $\text{Hist}(P) \triangleq \sum_{j \in [n]} \vec{v}_j$ to be the *histogram* of P , which represents the total weight of each combination of judgements in P . We define fractional votes and fractional profiles to present conditions in the theorem. Notice that agents’ votes are in \mathcal{V} , which means that they are non-fractional.

Quota Rules. A quota rule is a natural and well-studied generalization of proposition-wise majority rule with different thresholds. Formally, given any vector of *acceptance thresholds* (*threshold* in short), denoted by $\vec{q} = (q_1, \dots, q_{p+1}) \in [0, 1]^{p+1}$ and any vector of *tie-breaking criteria* (*breaking* in short), denoted by $\vec{d} = (d_1, \dots, d_{p+1}) \in \{0, 1\}^{p+1}$, we define the quota rule $r_{\vec{q}, \vec{d}}(P)$ as follows. For any profile $P \in (\mathcal{V}_{\text{frac}})^n$ and any $i \in [p + 1]$, we let $n_i = \sum_{\vec{\omega} \in \Omega_i} \text{Hist}(P)(\vec{\omega})$ denote the total weight of agents whose i -th judgement is 1, then apply the quota rule with threshold q_i and breaking d_i . That is,

$$\forall i \in [p + 1], r_{\vec{q}, \vec{d}}(P)(i) \triangleq \begin{cases} 1 & \text{if } n_i > q_i \cdot n \\ d_i & \text{if } n_i = q_i \cdot n \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where $r_{\vec{q}, \vec{d}}(P)(i)$ is the i -th component of $r_{\vec{q}, \vec{d}}(P)$. We say that a profile P is *tied* in ω_i if $n_i = q_i \cdot n$.

Doctrinal Paradox. A profile P is said to be a doctrinal paradox under r , if $r(P)$ is inconsistent with f . Formally, we have the following definition.

Definition 1 (Doctrinal paradox). *Given any $n \in \mathbb{N}$, any logical connection function f , and any quota rule r , a profile $P \in \mathcal{V}_{\text{frac}}^n$ is a doctrinal paradox, if $r(P)(p + 1) \neq f(r(P)(1), \dots, r(P)(p))$.*

The following example illustrates a fractional profile, a quota rule, and the aggregated judgements.

Example 2. *A (fractional) profile P of three votes, a quota rule r (which uses the majority rule with the breaking in favor of 1 on both premises and in favor of 0 on the conclusion), and the aggregation is shown in Table 2.*

	weight	ω_1	ω_2	ϕ	\vec{v}
Agent 1	1	1	0	0	$(0, 0, 1, 0)$
Agent 2	1	0	1	0	$(0, 1, 0, 0)$
Agent 3	0.5	1	0	0	$(0, 0, 0.5, 0.5)$
	0.5	1	1	1	
Hist(P)					$(0, 1, 1.5, 0.5)$
n_i		2	1.5	0.5	
Breaking \vec{d}		1	1	0	
Threshold \vec{q}		0.5	0.5	0.5	
$\vec{q} \cdot n$		1.5	1.5	1.5	
$r_{\vec{q}, \vec{d}}(P)$		1	1	0	

Table 2: The profile, rule, and aggregation result for Example 2.

It can be seen from the table that $r(P) = (1, 1, 0)$, which is inconsistent w.r.t. f . Therefore, P is a doctrinal paradox.

Characterization of Semi-random Likelihood of Doctrinal Paradox

Intuition. Before formally present the theorem, we first take a high-level and intuitive approach to introduce the idea behind the characterization and highlight the challenges. Take the max part $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ for example. Suppose the (max-)adversary has chosen $\vec{\pi} = (\pi_1, \dots, \pi_n) \in \Pi^n$. Let $\hat{\pi} = \sum_{j=1}^n \pi_j/n$ denote the center of $\vec{\pi}$. Intuitively, various multivariate central limit theorems imply that the histogram of the profile P generated from $\vec{\pi}$ falls in an $\Theta(n^{-0.5})$ neighborhood of $n\hat{\pi}$, denoted by $\mathcal{N}_{n\hat{\pi}}$, with high probability. As an approximation and relation, for now we allow the max-adversary to choose any $\hat{\pi}$ in the convex hull of Π (denoted by $\text{CH}(\Pi)$). Then, we can focus on the likelihood of doctrinal paradox in $\mathcal{N}_{n\hat{\pi}}$, which leads to the following four cases.

Case 1. If no profile of n votes is a doctrinal paradox, then $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n) = 0$ by definition.

Case 2. Otherwise, if for all $\hat{\pi} \in \text{CH}(\Pi)$, $\mathcal{N}_{n\hat{\pi}}$ does not contain a doctrinal paradox, then $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ is small. Notice that $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ is not zero, because there is a small (exponential) probability that $\text{Hist}(P)$ is not in $\mathcal{N}_{n\hat{\pi}}$.

Case 3. Otherwise, if the adversary can only choose $\hat{\pi} \in \text{CH}(\Pi)$ that corresponds to a “non-robust” instance of doctrinal paradox, in the sense that $\mathcal{N}_{n\hat{\pi}}$ includes some but not too many doctrinal paradoxes. Then, $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ should be larger than that in the previous case. This is the most challenging case, because it is unclear how to characterize conditions for it to happen, and how large the likelihood is. While standard techniques may lead to an $O(n^{-0.5})$ upper bound, it is unclear if $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ can be lower.

Case 4. Otherwise, the adversary can choose $\hat{\pi} \in \text{CH}(\Pi)$ that corresponds to a “robust” instance of doctrinal paradox, such that $\mathcal{N}_{n\hat{\pi}}$ contains many paradoxes. In this case, $\widetilde{\text{DP}}_{\Pi,r,f}^{\max}(n)$ is large.

Case 1 is trivial. It turns out that Case 2 and Case 3 can be characterized by properties of distributions in the convex hull of Π . We first define *effective refinements* of distributions to reason about aggregated judgements in $\mathcal{N}_{n\hat{\pi}}$.

Definition 2 ((Effective) Refinements). For any distribution π over \mathcal{V} , any $n \in \mathbb{N}$, and any quota rule r , we say that a vector $\vec{\alpha} = (\alpha_1, \dots, \alpha_{p+1}) \in \{0, 1\}^{p+1}$ is an effective refinement of π (at n), if the following two conditions hold:

- (1) $\forall i \in [p+1]$ such that π is not tied in ω_i , $\alpha_i = r(\pi)(i)$
- (2) $\exists P \in \mathcal{V}^n$ such that $r(P) = \vec{\alpha}$.

If condition (1) holds, then we call $\vec{\alpha}$ a refinement of π . Let \mathcal{E}_π denote the set of all effective refinements of π .

In words, condition (1) requires $\vec{\alpha}$ to match the aggregated judgements on all non-tied propositions. Condition (2) re-

quires the existence of a profile of n votes, whose outcome under r is $\vec{\alpha}$.

Example 3. In the setting of Example 2, $\pi = (0.3, 0.2, 0, 0.5)$ is tied in the first premise (ω_1) and the conclusion. There are four refinements of π : $(0, 1, 0)$, $(0, 1, 1)$, $(1, 1, 0)$ and $(1, 1, 1)$. $(0, 1, 1)$ is not effective, because no profile can make conclusion $\phi = 1$ while keeping ω_1 to be 0. The other three refinements are effective when $n \geq 3$. For example, $(1, 1, 1)$ is effective because it is the aggregation of the profile where all agents have 1 judgements on both premises.

Recall that $\text{CH}(\Pi)$ denotes the convex hull of Π . Next, we define four conditions (κ_1 to κ_4) to present Theorem 1. While their formal definitions may appear technical, we believe that they have intuitive explanations as discussed right afterward. In fact, κ_1 , κ_2 , and κ_4 correspond to case 1,2,3 discussed in the beginning of this section, respectively. κ_3 is defined for min-semi-random likelihood.

Definition 3 (Conditions κ_1 to κ_4). Given p premises, a logical connection function f , $n \in \mathbb{N}$, and a quota rule $r_{\vec{q}, \vec{d}}$ we define the following conditions.

κ_1 : $\forall P \in \mathcal{V}^n$, P is not a doctrinal paradox.

κ_2 : $\forall \pi \in \text{CH}(\Pi)$, $\forall \vec{\alpha} \in \mathcal{E}_\pi$, $\vec{\alpha}$ is consistent (under f).

κ_3 : $\exists \pi \in \text{CH}(\Pi)$ such that $\forall \vec{\alpha} \in \mathcal{E}_\pi$, $\vec{\alpha}$ is consistent (under f).

κ_4 : $\exists i \in [m]$ such that “ $\omega_{p+1} \leftrightarrow \omega_i$ and $q_{p+1} = q_i$ ” or “ $\omega_{p+1} \leftrightarrow \neg \omega_i$ and $q_{p+1} = 1 - q_i$ ”.

Intuitively, κ_1 requires that no profile of n votes is a doctrinal paradox. κ_2 requires that the output of $r_{\vec{q}, \vec{d}}$ of every distribution π in $\text{CH}(\Pi)$ is “far away” from any inconsistent judgements, in the sense that all judgements in \mathcal{E}_π are consistent. κ_3 is weaker than κ_2 , because κ_3 only requires that the condition in κ_2 holds for *some* $\pi \in \text{CH}(\Pi)$. κ_4 says that the conclusion only relies on one premise and the thresholds are consistent with the relationship between the conclusion and the premise.

Example 4 (Conditions κ_1 - κ_4). Let f and r be the same as in Example 2. Let $\Pi = \{\pi_1 = (0.25, 0.25, 0.25, 0.25), \pi_2 = (0.04, 0.32, 0.32, 0.32)\}$. When $n = 1$, the paradox does not hold by definition. In the rest of this example, we assume $n \geq 2$.

κ_1 is false. When $n = 2$, $P = (1, 0, 0, 1)$ is the only doctrinal paradox. When $n \geq 3$, $P = (n+1 - 2\lfloor n/3 \rfloor - \lceil n/3 \rceil, \lfloor n/3 \rfloor, \lfloor n/3 \rfloor, \lceil n/3 \rceil - 1)$ is a doctrinal paradox.

κ_2 is false. $\text{CH}(\Pi) = \{\pi = a \cdot \pi_1 + (1-a) \cdot \pi_2 : a \in [0, 1]\}$. When $a \neq 0$, π is a doctrinal paradox and contains no ties. When $a = 0$, we have $\pi = \pi_1$, and it is tied in conclusion and both premises. Its effective refinement $\vec{\alpha} = (1, 1, 0)$ is a doctrinal paradox.

κ_3 is false. π_2 is a doctrinal paradox and contains no ties.

κ_4 is false according to the definitions of f and r .

We say that a distribution π is *strictly positive*, if there exists $\epsilon > 0$ such that for every $\vec{\omega} \in \{0, 1\}^p$, $\pi(\vec{\omega}) \geq \epsilon$. A set of distributions Π is *strictly positive*, if there exists $\epsilon > 0$ such that every $\pi \in \Pi$ is strictly positive (by ϵ). Being strictly positive is a mild requirement for voting systems (c.f. the argument for voting (Xia 2020)). We also note that specifying

the likelihood of doctrinal paradox is already highly challenging under IC, which is a much stronger assumption than *strictly positive*. Π is *closed* if it is a closed set in \mathbb{R}^m . We now present the main theorem of this paper.

Theorem 1 (Semi-random Likelihood of Doctrinal Paradox). *Given any closed and strictly positive set Π of distributions over \mathcal{V} , any logical connection function $f : \mathcal{V} \rightarrow \{0, 1\}$, and any quota rule r . For any $n \in \mathbb{N}$,*

$$\widetilde{\text{DP}}_{\Pi, r, f}^{\max}(n) = \begin{cases} 0 & \text{if } \kappa_1 \text{ is true} \\ \exp(-\Theta(n)) & \text{otherwise, if } \kappa_2 \text{ is true} \\ \Theta(n^{-1/2}) & \text{otherwise, if } \kappa_4 \text{ is true} \\ \Theta(1) & \text{otherwise} \end{cases}$$

$$\widetilde{\text{DP}}_{\Pi, r, f}^{\min}(n) = \begin{cases} 0 & \text{if } \kappa_1 \text{ is true} \\ \exp(-\Theta(n)) & \text{otherwise, if } \kappa_3 \text{ is true} \\ \Theta(n^{-1/2}) & \text{otherwise, if } \kappa_4 \text{ is true} \\ \Theta(1) & \text{otherwise} \end{cases}$$

Notice that both max and min semi-random likelihood of doctrinal paradox have four cases: the 0 case, the exponential case, the polynomial case, and the constant case. The first three cases are good news, particularly in the max part, which states that the doctrinal paradox vanishes when the number of agents is large. The last case is bad news, in particular for the min part, which states that the doctrinal paradox does not vanish.

We believe that Theorem 1 is quite general as exemplified in the next subsection, because it works for any $p \geq 1$, any logical connection function, any quote rule, and only makes mild assumptions on Π . We believe that asymptotic studies as done in Theorem 1 is important, because first, n is large in some classical applications such as combinatorial voting (Lang and Xia 2016). Second, even though n is typically small in some classical applications such as jury systems, we believe that modern technology will revolutionize such applications in the near future and promote them to the larger scale and broader applications. In this case, Theorem 1 allows us to analyze the practical relevance of the paradox, and provides a mathematical basis for choosing the best (quote) rule to minimize the likelihood of the paradox.

Applications of Theorem 1

In this subsection, we present a few examples of applications of Theorem 1. The first example illustrates the $\Theta(1)$ case.

Example 5 ($\Theta(1)$ case). *Recall that in the setting of Example 4, we have: $(\kappa_1, \kappa_2, \kappa_3, \kappa_4) =$*

$$\begin{cases} (\text{true}, \text{false}, \text{false}, \text{false}) & \text{if } n = 1 \\ (\text{false}, \text{false}, \text{false}, \text{false}) & \text{if } n \geq 2 \end{cases}$$

According to Theorem 1, the max and min semi-random likelihood of doctrinal paradox are $\Theta(1)$, as shown in Table 3. A numerical verification can be found in Figure 2(a) in the next section.

The second example illustrates the exponential case, which is positive news.

Example 6 (Exponentially small case). *Let us consider the same quota rule and the same logical connection as in Example 2 and 4. Let $\Pi = \{\pi_1 =$*

	$n = 1$	$n \geq 2$ (Ex. 5)	$n \geq 2$ (Ex. 6)
$\widetilde{\text{DP}}_{\Pi, r_{\vec{q}, \vec{d}}, f}^{\max}(n)$	0	$\Theta(1)$	$\exp(-\Theta(n))$
$\widetilde{\text{DP}}_{\Pi, r_{\vec{q}, \vec{d}}, f}^{\min}(n)$	0	$\Theta(1)$	$\exp(-\Theta(n))$

Table 3: Semi-random likelihood of DP in Example 5 and 6.

$(0.12, 0.12, 0.12, 0.64), \pi_2 = (0.1, 0.1, 0.1, 0.7)$. Following a similar reasoning as in Example 4, we have:

$$(\kappa_1, \kappa_2, \kappa_3, \kappa_4) = \begin{cases} (\text{true}, \text{false}, \text{false}, \text{false}) & \text{if } n = 1 \\ (\text{false}, \text{true}, \text{true}, \text{false}) & \text{if } n \geq 2 \end{cases}$$

According to Theorem 1, the max and min semi-random likelihood of doctrinal paradox are $\exp(-\Theta(n))$, as shown in Table 3. See Figure 2(b) for a numerical verification.

The following example is also positive news because the doctrinal paradox vanishes as $n \rightarrow \infty$, though not as fast as in Example 6 w.r.t. the max-adversary. Notice that the max and min semi-random likelihood of doctrinal paradox are asymptotically different.

Example 7 ($\Theta(n^{-1/2})$, $\exp(-\Theta(n))$ and 0 cases).

Let the logical connection be $\phi \leftrightarrow \omega_1$, $\Pi = \{\pi_1 = (0.9, 0.1), \pi_2 = (0.3, 0.7)\}^1$ and quota rule $r_{\vec{q}, \vec{d}}$ where $q_1 = q_3 = 0.5$, $d_1 = 1$, and $d_3 = 0$. We have: $(\kappa_1, \kappa_2, \kappa_3, \kappa_4) =$

$$\begin{cases} (\text{true}, \text{false}, \text{false}, \text{true}) & \text{if } n \text{ is odd} \\ (\text{false}, \text{false}, \text{true}, \text{true}) & \text{if } n \text{ is even} \end{cases}$$

According to Theorem 1, the max and min semi-random likelihood are presented in Table 4. Also see Figure 2(c) the next Section for a numerical verification.

Semi-random likelihood	n is odd	n is even
$\widetilde{\text{DP}}_{\Pi, r_{\vec{q}, \vec{d}}, f}^{\max}(n)$	0	$\Theta(n^{-1/2})$
$\widetilde{\text{DP}}_{\Pi, r_{\vec{q}, \vec{d}}, f}^{\min}(n)$	0	$\exp(-\Theta(n))$

Table 4: Semi-random likelihood of DP in Example 7.

The following corollary of Theorem 1 with $\Pi = \{\pi\}$ address the open questions in (List 2005) about the probabilities of doctrinal paradox under all i.i.d. distributions, all q , all quota rules, and all logical connection functions.

Corollary 2 (Likelihood of doctrinal paradox under i.i.d. distributions as in (List 2005)). *Given any strictly positive distribution π over \mathcal{V} , any logical connection function $f : \mathcal{V} \rightarrow \{0, 1\}$, and any quota rule r . For any $n \in \mathbb{N}$,*

$$\Pr_{P \sim (\pi)^n} (P \text{ is a doctrinal paradox}) =$$

$$\begin{cases} 0 & \text{if } \kappa_1 \text{ is true} \\ \exp(-\Theta(n)) & \text{otherwise, if all } \vec{\alpha} \in \mathcal{E}_\pi \text{ are consistent,} \\ \Theta(n^{-1/2}) & \text{otherwise, if } \kappa_4 \text{ is true} \\ \Theta(1) & \text{otherwise} \end{cases}$$

¹Because the doctrinal paradox only depends on the votes to premise ω_j (or conclusion ϕ), we use the marginal distribution on ω_i to simplify notations. Here, $\pi = (\text{pr}, 1 - \text{pr})$ mean $\omega_i = 0$ with probability pr while $\omega_i = 1$ with probability $1 - \text{pr}$.

In particular, when π is the uniform distribution over \mathcal{V} and the aggregation rule is the majority, the likelihood of doctrinal paradox is either $\Theta(1)$ or 0 depending on the logical connection function f .

Proof Sketch of Theorem 1

The proof proceeds in three steps. In **Step 1**, we model doctrinal paradox as unions of polyhedra. In **Step 2**, we apply (Xia 2021a, Theorem 2) to obtain a characterization. In **Step 3**, we close the gap between Step 2 and Theorem 1 by translating conditions and degree of polynomial in Step 2 to their counterparts in Theorem 1.

Step 1: Polyhedra Presentation. We show that the region of doctrinal paradox can be presented by a set of polyhedra, each of which is in the m -dimensional Euclidean space with the form $\mathcal{H} = \{\vec{x} : \mathbf{A} \cdot \vec{x} \leq \vec{b}\}$. For any $i \in [p+1]$, we first define vector $\vec{c}_i \in \{q_i - 1, q_i\}^m$ as follows:

$$\vec{c}_i(\vec{\omega}) \triangleq \begin{cases} q_i - 1 & \text{if } \vec{\omega} \in \Omega_i \\ q_i & \text{otherwise} \end{cases} \quad (2)$$

For any profile P , it is not hard to verify that the sign of $\vec{c}_i \cdot \text{Hist}(P)$ is closely related to the aggregation result by the quota rule. We define sign function $\text{sign}(\cdot)$ and $1(\cdot)$ as follows,

$$\text{sign}(x) \triangleq \begin{cases} 1 & \text{if } x > 0 \\ -1 & \text{otherwise} \end{cases} \quad \text{and} \\ 1(\kappa) \triangleq \begin{cases} 1 & \text{if } \kappa \text{ is true} \\ 0 & \text{otherwise} \end{cases}$$

Then, we define $\mathbf{A}_{\vec{\alpha}} \triangleq \begin{pmatrix} \text{sign}(\alpha_1) \cdot \vec{c}_1^\top \\ \vdots \\ \text{sign}(\alpha_{p+1}) \cdot \vec{c}_{p+1}^\top \end{pmatrix}$ and $\vec{b}_{\vec{\alpha}} \triangleq \begin{pmatrix} \text{sign}(\alpha_1) \cdot d_1 - 1(\alpha_1 = 1) \\ \vdots \\ \text{sign}(\alpha_{p+1}) \cdot d_{p+1} - 1(\alpha_{p+1} = 1) \end{pmatrix}$. Let $\mathcal{H}_{\vec{\alpha}} = \{\vec{x} :$

$\mathbf{A}_{\vec{\alpha}} \vec{x} \leq \vec{b}_{\vec{\alpha}}\}$ denote the polyhedra and let $\mathcal{H}_{\vec{\alpha}, \leq 0} = \{\vec{x} : \mathbf{A}_{\vec{\alpha}} \vec{x} \leq \vec{0}\}$ denote its *characteristic cone*. Let $\mathcal{C} = \{\mathcal{H}_{\vec{\alpha}} : \vec{\alpha} \in \{0, 1\}^{p+1} \text{ is inconsistent}\}$ denote the set of all polyhedra corresponding to doctrinal paradox. The following example visualizes a polyhedron and its characteristic cone for doctrinal paradox.

Example 8 (Polyhedra representation of doctrinal paradox). We consider a system with one premise ω_1 . The logical connection between conclusion ϕ and premise ω_1 is $\phi \leftrightarrow \omega_1$. The parameters for quota rule are $(q_1, q_2) = (0.25, 0.65)$ and $(d_1, d_2) = (1, 1)$.

According to Equation (2), we have $\vec{c}_1 = (q_1, q_1 - 1)^\top$ and $\vec{c}_2 = (q_2, q_2 - 1)^\top$. Then, the aggregation results of $\vec{\alpha} = (1, 0)$ and $\vec{\alpha} = (0, 1)$ are both inconsistent. Thus, $\mathcal{C} = \{\mathcal{H}_{(1,0)}, \mathcal{H}_{(0,1)}\}$. $\mathcal{H}_{(1,0)}$ is represented by $\mathbf{A}_{(1,0)}$ and $\vec{b}_{(1,0)}$ defined as follows.

$$\mathbf{A}_{(1,0)} = \begin{pmatrix} q_1 & q_1 - 1 \\ -q_2 & 1 - q_2 \end{pmatrix} \quad \text{and} \quad \vec{b}_{(1,0)} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$




Figure 1: $\mathcal{H}_{(1,0)}$ in Example 8. The characteristic cone $\mathcal{H}_{(1,0), \leq 0}$ is the combination of $\mathcal{H}_{(1,0)}$ and the region between $n_1 = q_2 \cdot n - 1$ and $n_1 = q_2 \cdot n$.

It is not hard to verify that for any profile P of n agents, $\mathbf{A}_{(1,0)} \cdot \text{Hist}(P) \leq \vec{b}_{(1,0)}$ if and only if

$$n_1 \geq q_1 \cdot n \quad \text{and} \quad n_1 \leq q_2 \cdot n - 1,$$

where n_1 represents the number of votes for $\omega_1 = 1$ in P . Figure 1 illustrates the regions corresponding to $\mathcal{H}_{(1,0)}$ and its characteristic cone.

Consequently, the semi-random likelihood of doctrinal paradox is equivalent to the semi-random probability for the histogram of the randomly generated profile (which is a *Poisson multivariate variable (PMV)*) to be in \mathcal{C} . Therefore, in Step 2 we apply (Xia 2021a, Theorem 2) to obtain a characterization for the PMV-in- \mathcal{C} problem. However, the results are too technical and generic to be informative in judgement aggregation context (see Lemma 5 in Appendix B). Step 3 aims at obtaining the characterization as stated in Theorem 1 based on Lemma 5, which involves non-trivial calculations and simplifications. The full proof can be found in Appendix B. \square

Computational Verification of Conditions in Theorem 1

In practice, $\kappa_1 - \kappa_4$ may not be easy to verify by hand. In this subsection, we briefly discuss algorithms for checking $\kappa_1 - \kappa_4$ under any quota rule $r_{\vec{q}, \vec{d}}$ with rational threshold \vec{q} w.r.t. any Π with finitely many vertices on its convex hull and the probabilities of all vertices are rational numbers. The runtime of these algorithms is polynomial in m (as well as $\log n$ for κ_1). While $m = 2^p$ may seem exponentially large, notice that the input size (e.g., Π and the truth-table representation of f) can also be large, i.e., $\Theta(m)$.

Suppose $\Pi = \{\pi^1, \dots, \pi^\ell\} \subseteq \mathbb{Q}^{m, 2}$. $\kappa_1 - \kappa_4$ can be verified as follow, by combining the results of multiple (I)LP

²Throughout this section, we slightly abuse the notations and use Π to represent the set of all vertices in $\text{CH}(\Pi)$.

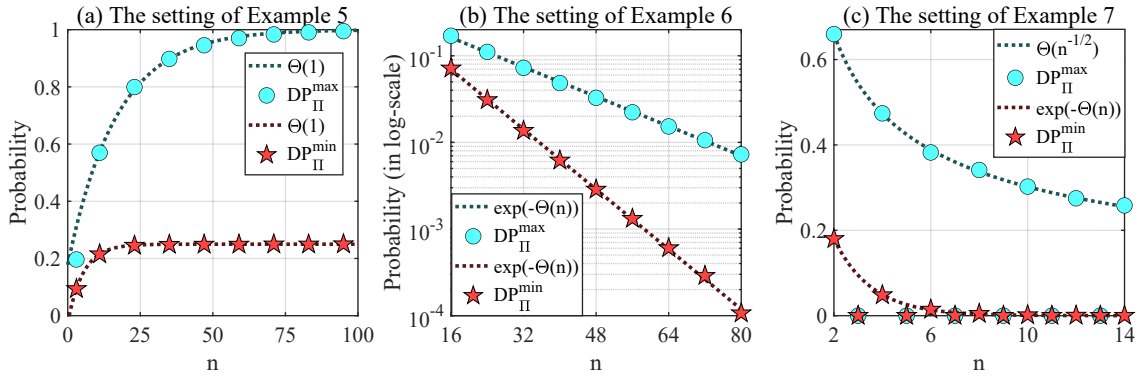


Figure 2: A numerical verification of Theorem 1. Note that the plot for Example 6 is in log-scale.

whose variables are $\vec{y} = (y^1, \dots, y^\ell)$ that represent the multiplicity of π^ℓ .

κ_1 holds for any given n if and only if there exists an inconsistent $\vec{\alpha} \in \{0, 1\}^{p+1}$ such that the following ILP is feasible.

$$\begin{pmatrix} \mathbf{A}_{\vec{\alpha}} \cdot ((\pi^1)^\top, \dots, (\pi^\ell)^\top) \\ \vec{1} \\ -\vec{1} \end{pmatrix} \cdot (\vec{y}) \leq \begin{pmatrix} \vec{b}_{\vec{\alpha}} \\ n \\ -n \end{pmatrix}, \quad (3)$$

where \vec{y} are integers. Following Lenstra's theorem (Lenstra 1983), when ℓ is viewed as a constant (and q and n are variables), the runtime of each ILP is polynomial in q and $\log n$. Therefore, the overall complexity of checking κ_1 is polynomial in m ($= 2^p$) and $\log n$. Let $A_n \subseteq \{0, 1\}^{p+1}$ denote the set of all inconsistent α such that (3) has a feasible (non-negative integer) solution, which will be used later.

κ_2 and κ_3 can be verified by similar algorithms. Notice that \mathcal{E}_π only depends on whether each component of $r(\pi)$ is positive, 0, or negative. Therefore, we can naturally define $\mathcal{E}_{\vec{\beta}}$ for every $\vec{\beta} \in \{+, 0, -\}^{p+1}$. Let B_n denote the set of all $\vec{\beta}$, for each of which there exists $\alpha \in A_n$ that refines it. Then, for every $\vec{\beta} \in B_n$, we define $\text{LP}_{\vec{\beta}}$, whose variables are $\vec{y} \geq \vec{0}$ that do not need to be integers, to verify whether there exists $\pi \in \text{CH}(\Pi)$ such that $r_{\vec{q}, \vec{d}}(\pi) = \vec{\beta}$. More precisely, for every $j \leq p+1$,

- 1° if $\beta_j = +$, then $\text{LP}_{\vec{\beta}}$ contains a linear constraint $\vec{c}_j^\top \cdot (\sum_{j=1}^\ell y_j (\vec{\pi}^j)^\top) \leq -1$;
- 2° if $\beta_j = -$, then $\text{LP}_{\vec{\beta}}$ contains a linear constraint $-\vec{c}_j^\top \cdot (\sum_{j=1}^\ell y_j (\vec{\pi}^j)^\top) \leq -1$;
- 3° if $\beta_j = 0$, then $\text{LP}_{\vec{\beta}}$ contains two linear constraints $\vec{c}_j^\top \cdot (\sum_{j=1}^\ell y_j (\vec{\pi}^j)^\top) \leq 0$ and $-\vec{c}_j^\top \cdot (\sum_{j=1}^\ell y_j (\vec{\pi}^j)^\top) \leq 0$.

Therefore, κ_2 is true if for every $\vec{\beta} \in B_n$, $\text{LP}_{\vec{\beta}}$ is not feasible. κ_3 is true, if there exists $\vec{\beta} \notin B_n$ such that $\text{LP}_{\vec{\beta}}$ is feasible. Notice that ℓ is viewed as a constant, which means that $\text{LP}_{\vec{\beta}}$ can be solved in time that is polynomial in m .

κ_4 can be verified by enumerating all inputs of f , which takes time that is polynomial in m .

Experiments

We conduct numerical experiments to verify the results in Theorem 1. The first three experiments (Figure 2) follows the same setting as Example 5, 6, 7 respectively. In Figure 2, DP_{Π}^{\max} and DP_{Π}^{\min} (the blue circles and red stars) represent the estimated max-semi-random likelihood and the min-semi-random likelihood of doctrinal paradoxes. The dot curves illustrate the fittings of the estimated semi-random probabilities. The expressions and fitness of all fitting curves are presented in Appendix C. Recalling the notations used in our definition of $\text{DP}_{\Pi, r, f}^{\max}(n)$, we say $\vec{\pi} \in \Pi^n$ is one kind of *distribution assignment*. We run one million (10^6) independent trials to estimate the probability of doctrinal paradoxes under each distribution assignment. Then, DP_{Π}^{\max} (or DP_{Π}^{\min}) takes the maximum (or the minimum) probability of doctrinal paradoxes among all distribution assignments.

It is easy to see that the results are consistent with Theorem 1. For example, Figure 2(a) shows that the max-semi-random likelihood and the min-semi-random likelihood of doctrinal paradoxes are $\Theta(1)$, which matches the result in Table 3. Two additional sets of experiments, whose results are presented in Figure 3(b) in Appendix C, study a more complex setting of three premises.

Conclusions and Future Work

We made a first step towards understanding the likelihood of doctrinal paradox in the semi-random analysis framework. There are many immediate open questions. For example, can we characterize the semi-random likelihood of the paradox for general, more complicated propositions (sometimes called *agendas*)? Can we remove the strict positive-ness assumption on Π ? Can we prove semi-random impossibility theorems that extend the quantitative impossibility theorems under i.i.d. uniform distributions (Nehama 2013; Filmus et al. 2020)? Can we generalize the framework to discursive dilemmas or other kinds of responsibility voids? More generally, we believe that building a comprehensive picture of the semi-random properties of aggregation rules in judgement aggregation is a promising (yet challenging) direction in theory and in practice.

Acknowledgements

We thank anonymous reviewers and COMSOC-21 attendees for helpful comments. Lirong Xia is supported by NSF #1453542, ONR #N00014-17-1-2621, and a gift fund from Google. Ao Liu is supported by an RPI-IBM AI Horizons scholarship.

References

- Awad, E.; Booth, R.; Tohmé, F.; and Rahwan, I. 2017. Judgement aggregation in multi-agent argumentation. *Journal of Logic and Computation*, 27(1): 227–259.
- Baharad, E.; Neeman, Z.; and Rubinchik, A. 2020. The rarity of consistent aggregators. *Mathematical Social Sciences*, 108: 146–149.
- Baumeister, D.; Erdelyi, G.; Erdelyi, O. J.; Rothe, J.; and Selker, A.-K. 2020. Complexity of control in judgment aggregation for uniform premise-based quota rules. *Journal of Computer and System Sciences*, 112: 13–33.
- Blum, A.; and Spencer, J. 1995. Coloring random and semi-random k -colorable graphs. *Journal of Algorithms*, 19(2): 204–234.
- Bonnefon, J.-F. 2010. Behavioral evidence for framing effects in the resolution of the doctrinal paradox. *Social Choice and Welfare*, 34: 631–641.
- Bonnefon, J.-F. 2011. The doctrinal paradox, a new challenge for behavioral psychologists. *Advances in Psychological Science*, 19(5): 617–623.
- Braham, M.; and VanHees, M. 2011. Responsibility voids. *The Philosophical Quarterly*, 61(242): 6–15.
- Brandl, F.; Brandt, F.; and Seedig, H. G. 2016. Consistent probabilistic social choice. *Econometrica*, 84(5): 1839–1880.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Chilton, A.; and Tingley, D. 2012. The doctrinal paradox & international law. *University of Pennsylvania Journal of International Law*, 34(67).
- Dietrich, F.; and List, C. 2008. Judgment aggregation without full rationality. *Social Choice and Welfare*, 31(1): 15–39.
- Diss, M.; and Merlin, V., eds. 2021. *Evaluating Voting Systems with Probability Models*. Studies in Choice and Welfare. Springer International Publishing.
- Endriss, U. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press.
- Endriss, U.; de Haan, R.; Lang, J.; and Slavkovik, M. 2019. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of Artificial Intelligence Research*, 69: 687–731.
- Endriss, U.; Grandi, U.; and Porello, D. 2012. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research*, 45: 481–514.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2017. An Introduction to Belief Merging and its Links with Judgment Aggregation. In *Trends in Computational Social Choice*, chapter 7, 123–143. AI Access.
- Filmus, Y.; Lifshitz, N.; Minzer, D.; and Mossel, E. 2020. AND testing and robust judgement aggregation. In *Proceedings of STOC*.
- Gehrlein, W. V.; and Lepelley, D. 2011. *Voting Paradoxes and Group Coherence: The Condorcet Efficiency of Voting Rules*. Springer.
- Grossi, D.; and Pigozzi, G. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool Publishers.
- Hanna, C. 2009. The paradox of progress: Translating even stark’s coercive control into legal doctrine for abused women. *Violence Against Women*, 15(2): 1458–1476.
- Kornhauser, L. A.; and Sager, L. G. 1986. Unpacking the Court. *The Yale Law Journal*, 96(1): 82–117.
- Lang, J.; and Xia, L. 2016. Voting in Combinatorial Domains. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press.
- Lehtinen, A.; and Kuorikoski, J. 2007. Unrealistic assumptions in rational choice theory. *Philosophy of the Social Sciences*, 37(2): 115–138.
- Lenstra, H. W. 1983. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8: 538–548.
- List, C. 2005. The probability of inconsistencies in complex collective decisions. *Social Choice and Welfare*, 24: 3–32.
- List, C. 2012. The theory of judgment aggregation: an introductory review. *Synthese*, 187: 179–207.
- List, C.; and Pettit, P. 2002. Aggregating sets of judgments: An impossibility result. *Economics and philosophy*, 18(1): 89–110.
- List, C.; and Polak, B. 2010. Introduction to judgment aggregation. *Journal of Economic Theory*, 145(2): 441–466.
- Lyon, A.; and Pacuit, E. 2013. The wisdom of crowds: Methods of human judgement aggregation. In *Handbook of human computation*, 599–614. Springer.
- Marcoci, A.; and Nguyen, J. 2020. Judgement aggregation in scientific collaborations: The case for waiving expertise. *Studies in History and Philosophy of Science Part A*, 84: 66–74.
- Mongin, P. 2008. Factoring out the impossibility of logical aggregation. *Journal of Economic Theory*, 141(1): 100–113.
- Mongin, P. 2012. The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory and decision*, 73(3): 315–355.
- Mongin, P. 2019. The Present and Future of Judgement Aggregation Theory. A Law and Economics Perspective. In *The future of economic design*, 417–429. Springer.
- Nehama, I. 2011. Approximate judgement aggregation. In *International Workshop on Internet and Network Economics*, 302–313. Springer.

- Nehama, I. 2013. Approximately classic judgement aggregation. *Annals of Mathematics and Artificial Intelligence*, 68: 91–134.
- Nehring, K.; and Pivato, M. 2021. The median rule in judgement aggregation. *Economic Theory*.
- Nehring, K.; and Puppe, C. 2008. Consistent judgement aggregation: the truth-functional case. *Social Choice and Welfare*, 31(1): 41–57.
- Nurmi, H. 1999. *Voting Paradoxes and How to Deal with Them*. Springer-Verlag Berlin Heidelberg.
- Pauly, M.; and Hees, M. V. 2006. Logical constraints on judgement aggregation. *Journal of Philosophical Logic*, 35(6): 569–585.
- Pettit, P. 2001. Deliberative democracy and the discursive dilemma. *Philosophical Issues*, 11: 268–299.
- Poisson, S. D. 1837. *Recherches sur la probabilité des jugements en matière criminelle et en matière civile*.
- Rahwan, I.; and Tohmé, F. 2010. Collective argument evaluation as judgement aggregation. In *Proceedings of AAMAS*.
- Sorensen, R. 2003. *A Brief History of the Paradox: Philosophy and the Labyrinths of the Mind*. Oxford University Press.
- Spielman, D. A.; and Teng, S.-H. 2009. Smoothed Analysis: An Attempt to Explain the Behavior of Algorithms in Practice. *Communications of the ACM*, 52(10): 76–84.
- Vacca, R. 1921. Opinioni individuali e deliberazioni collettive. *Rivista Internazionale di Filosofia del Diritto*, 52(52–59).
- Van Deemen, A. 2014. On the empirical relevance of Condorcet’s paradox. *Public Choice*, 158(3): 311–330.
- Xia, L. 2020. The Smoothed Possibility of Social Choice. In *Proceedings of NeurIPS*.
- Xia, L. 2021a. How Likely Are Large Elections Tied? In *Proceedings of ACM EC*.
- Xia, L. 2021b. The Semi-Random Satisfaction of Voting Axioms. *Advances in Neural Information Processing Systems*, 34.
- Xia, L.; and Zheng, W. 2021. The Smoothed Complexity of Computing Kemeny and Slater Rankings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5742–5750.
- Zhang, H.; and Conitzer, V. 2019. A PAC Framework for Aggregating Agents’ Judgments. In *Proceedings of AAAI*.