Mathematical Social Sciences 108 (2020) 146–149

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/mss

The rarity of consistent aggregators[☆]

Eyal Baharad[a,*], Zvika Neeman[b], Anna Rubinchik[c]

[a] Department of Economics; Bar Ilan University, Israel
[b] Department of Economics; Tel Aviv University, Israel
[c] Department of Economics; University of Haifa, Israel

ARTICLE INFO

Article history:
Received 11 January 2017
Received in revised form 7 January 2019
Accepted 12 September 2019
Available online 28 May 2020

Keywords:
Aggregation of opinions
Doctrinal paradox

ABSTRACT

We demonstrate that the inconsistency associated with judgment aggregation, known as the "doctrinal paradox", is not a rare exception. There are n individuals who have opinions about k propositions. Each opinion expresses the degree of belief or conviction and thus belongs to the unit interval [0, 1]. We work with an arbitrary proposition aggregator that maps opinions about k propositions into an overall opinion in [0, 1] and an arbitrary individual opinions aggregator mapping opinions of n individuals into a single judgement from a unit interval. We show that for any typical proposition aggregator, the set of individual opinion aggregators that are immune to the paradox is very small, i.e., is nowhere dense in the space of uniformly bounded functions. In addition, we offer several examples of judgement aggregation for which the paradox can be avoided.

1. Introduction

The aggregation of binary judgments is well known to be problematic. We focus on a subclass of problems related to the so-called *doctrinal paradox* (cf. Kornhauser and Sager, 1993). We demonstrate the prevalence of this paradox in a general class of judgement aggregators.

Before formulating our result more precisely, we start with a brief overview of the related literature.

The classical description of the paradox proceeds as follows. Suppose that a three-judge court has to make a decision on whether a defendant is liable for breach of contract. According to the existing legal regime, the defendant is liable if and only if the contract is valid (proposition 1), and the contract was breached (proposition 2). Assume that the first judge is convinced that both propositions are true, the second judge believes that the contract was valid but it was not breached, and the third one believes the opposite, i.e., that the first proposition is false while the second is true. Thus, the matrix of opinions is

$$\begin{array}{cc|c} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array}$$

where the rows list the opinions of the three judges, the first two columns correspond to the propositions and the third one — to the conclusion of each judge. "Proposition-wise" majority voting would result in the assertion that both propositions should be true (2 against 1), thus in this case the defendant would be considered liable. However, only the first judge should conclude that the defendant is liable since he is the only one who believes that both propositions hold, and hence majority voting over the final decision yields the opposite verdict.

Note that under the first method for aggregating the judgements the initial step was to use majority voting to aggregate across individuals to form common *premises* which then were to be aggregated into a conclusion. This is a *premise*- or *reason*-based procedure. The second method of aggregation describes the *conclusion*- or *outcome*-based procedure.

Several contributions assess the relative merits of these two aggregation procedures for given aggregation rules (see, e.g., Bovens and Rabinowicz, 2004, 2006; Pettit, 2001). De Clippel and Eliaz (2015) consider an environment where every voter receives a signal regarding truth values of the premises and all voters agree on how to reach the final decision based on the actual premises that describe the state of nature. A super-majority rule aggregates the individual reports about either premises or individual decisions. The authors find the premise-based approach to be superior to the outcome-based in the following sense. The set of symmetric Bayesian Nash equilibria (SBNE) of the premise-based voters' game is a proper superset of the SBNE of the outcome-based game. Moreover in the presence of a separating set of signals distinguishing the two states, the SBNE in the premise-based game yields the efficient information aggregation: the probability of a discrepancy between the outcome of the game and the decision based on the actual state of nature is converging to zero as the number of voters approaches infinity.

☆ We would like to thank Prof. A. Gorokhovsky for insightful discussions.
* Corresponding author.
E-mail addresses: Eyal.Baharad@biu.ac.il (E. Baharad), zvika@post.tau.ac.il (Z. Neeman), arubinchik@econ.haifa.ac.il (A. Rubinchik).

https://doi.org/10.1016/j.mathsocsci.2019.09.007
0165-4896/

The next natural question is whether there is a way to avoid the paradox altogether, i.e., if we do not insist on majority aggregators, is there a way to reach a common conclusion that is independent of the procedure (premise- or conclusion-based)?

Dietrich and Mongin (2010) examine the implication of two axioms: independence on premises and unanimity preservation on premises and non-premises. They show when aggregation functions that satisfy these axioms are a dictatorship or an oligarchy.

By and large, if individual opinions are restricted to be binary, the answer to this question is negative (Dietrich, 2006; List and Pettit, 2002; List, 2004; List and Pettit, 2004; Dietrich, 2007; Nehring and Puppe, 2008), and the analysis has intricate connections to the well-known impossibility results of K. Arrow and A. Sen, see Dietrich and List (2007, 2008) and Dokow and Holzman (2010).

There is, however, an important caveat. The problem is typically formulated for a given set of rules that aggregate basic premises into conclusions or into *dependent (derived)* propositions. Thus aggregation rules across propositions have to abide by the laws of the underlying logic, or be *truth-functionally coherent* (Dietrich and List, 2010). With such a collection of consistency requirements it is often impossible to find a non-dictatorial rule that aggregates individual judgements, both basic and derived. This is true even if the individual judgements are restricted to be truth-functionally coherent themselves. The impossibility results extend even to multi-valued opinions (Guilbaud, 1952; Pauly and Van Hees, 2006).

Furthermore, it is possible to extend the domain to contain opinions that can describe the full range of judges' degrees of confidence, or *probability judgments*, i.e., allow the entries in the matrix of opinions to be from the unit interval [0, 1], see, for example, Nehring (2007) and List (2005) and a survey by Genest and Zidek (1986). We also follow this convention.

Dietrich and List (2010) offer a unified approach to the aggregation problem. With the opinions being represented by real numbers, basic premises have to be combined in a consistent way given either the language syntax or the standard probability calculus. The basic lesson (most closely related to the current study) is that under a variety of natural restrictions imposed by language or underlying uncertainty, the linear aggregation rule is the only one that works in a consistent way (Dietrich and List, 2017a,b). Herzberg (2013) proposes using multi-valued algebra as a framework for propositional-attitude aggregation: algebra homomorphism appears to be the only *desirable* aggregator of individual *attitude functions* that, naturally, satisfies the consistency requirements imposed by the language expressed as an algebraic structure. Thus, classical desiderata in this context are closely connected to linearity (see, also McConway, 1981).

We address the main question in an indirect way. First, we are agnostic about both the events and the context that give rise to the opinions of the judges. We start with the basic premises, much like the first two columns in the matrix of opinions above. We then work with two aggregators. The first one describes the relation between the basic premises and the conclusion, both on the individual and on the aggregate level (it generates the third column in the matrix of the example). The second one is the aggregator across individuals, which is most familiar to us from the social choice literature mentioned above.

In other words, our "language" imposes only a single validity restriction and this restriction is endogenously determined as the choice of one of the aggregators a-priori is not constrained. Given the first one, the output of the second aggregator is required to be consistent so as to avoid the doctrinal paradox. In the view of the results mentioned above it is not surprising that a pair of linear aggregators satisfy this restriction.

As we demonstrate in the next section, the impossibility results mentioned above cannot be extended to our setting even if the opinions are binary. One might conclude that in this case the paradox can be easily avoided. However, as we show, this conclusion is incorrect.

Our main result is that consistent aggregation is fragile on an unrestricted domain of opinions in the following sense. Choose a function that aggregates basic propositions into conclusions, or a deliberation rule. Then, for a given matrix of opinions, the set of all functions that can be used to aggregate across individuals while avoiding the paradox is very small, i.e., nowhere dense in the set of uniformly bounded functions.

2. Setup

Let $N = \{1, \dots, n\}$ denote a (finite) set of agents, and $K = \{1, \dots, k\}$ denote a (finite) set of *propositions*. Each agent $i \in N$ has an opinion $x_{ip} \in [0, 1]$ about each proposition $p \in K$ that expresses the agent's degree of self-persuasion about the truth value of the proposition. Let $X \in [0, 1]^{nk}$ be the matrix of opinions with elements $x_{ip} \in [0, 1]$, columns $y_i \in [0, 1]^k$ and rows $z_p \in [0, 1]^n$, $i \in N$, $p \in K$.

The objective is to aggregate the profile of agents' opinions, X, about propositions into a joint aggregate opinion: $[0, 1]^{nk} \to [0, 1]$.

A *proposition aggregator* function $f : [0, 1]^k \to [0, 1]$ is a mapping from a vector of opinions about k propositions into an opinion about the issue and an *agent aggregator* function $g : [0, 1]^n \to [0, 1]$ is a mapping from a vector of n opinions about a single proposition into an aggregate opinion. The two functions f and g can be combined to produce a joint aggregate opinion about the issue in either one of the following two ways.

- **First f then g**: $g \circ f^n$. First, aggregate across propositions to get individual opinions about the issue and then aggregate the latter into a joint opinion, $g\left(f(y_1), \dots, f(y_n)\right)$, $y_i \in [0, 1]^k$, $i \in N$.
- **First g then f**: $f \circ g^k$. First, aggregate across agents to get a joint opinion about each proposition, and then aggregate the latter into a common judgement about the issue $f\left(g(z_1), \dots, g(z_k)\right)$, $z_p \in [0, 1]^n$, $p \in K$.

If the joint aggregate opinion about the issue is independent of the order in which the two functions f and g are combined, so that $f \circ g^k = g \circ f^n$, then we say that *$f$ and g aggregate* on X. If g and f aggregate on every X, then we say they *aggregate consistently*. In this case the doctrinal paradox is avoided.

3. Possibility of aggregation

We first present a sufficient condition for aggregation.

Lemma 3.1. *If both f and g are linear functions, then they aggregate consistently.*

Proof. Let Δ^{n-1} be an $n-1$ simplex: $\Delta^{n-1} = \left\{v \in \mathbb{R}^n_+ | \sum_{i=1}^n v_i = 1\right\}$.

By linearity of f and g, $\exists\, \alpha \in \Delta^{k-1}$ and $\beta \in \Delta^{n-1}$ such that

$$f(y_i) = \alpha y_i = \sum_{p=1}^{k} \alpha_p x_{ip}, \quad \forall i \in N,$$

$$g(z_p) = \beta z_p = \sum_{i=1}^{n} \beta_i x_{ip} \quad \forall p \in K$$

Hence,

$$g(f(y_1),\dots,f(y_n)) = \sum_{i=1}^{n}\beta_i\sum_{p=1}^{k}\alpha_p x_{ip}$$
$$= \sum_{p=1}^{k}\alpha_p\sum_{i=1}^{n}\beta_i x_{ip} = f(g(z_1),\dots,g(z_k)) \quad \square$$

Example 1. Modulo of the sum aggregators of the form

$$y \mapsto \mod_a\left(\sum_{p=1}^{m} y_p\right), \quad y \in [0,1]^m, \quad 0 < a \le 1$$

aggregate consistently, i.e., if both the agent and the proposition aggregator are modulo of the sum, any order of aggregation leads to the same result. The same is true for the modulo of the product.

The reason for that is the distributive property of the modulo operation:

$$(c+b) \mod_a = [(c \mod_a) + (b \mod_a)] \mod_a$$
$$cb \mod_a = [(c \mod_a)(b \mod_a)] \mod_a$$

Example 2. If $n = k$ and the matrix X of opinions is symmetric then for any function h: $[0,1]^n \to [0,1]$, the pair f, g: $f = g = h$ aggregate consistently.

4. The rarity result

Note that all of our aggregators are bounded functions since their range is limited to the unit interval and therefore they form a subset Ω of Borel functions in $L_\infty([0,1]^n)$, or L_∞ for short.[1]

That dictatorial functions can aggregate is well-known, so we assume them away. The definition of a dictatorial function is standard.

Definition 1. A function h: $\mathbb{R}^m \to \mathbb{R}$ is dictatorial if it is a projection, i.e., there is $i \in \{1,\dots,m\}$ such that $h(x_1,\dots,x_m) = x_i$ for all $x \in \mathbb{R}^m$.

Our main statement in this section shows that being able to find f that aggregates consistently with any given simple function g is "hard".

First we offer the underlying geometric intuition for the argument for $n = k = 2$. Take a non-dictatorial function $g \in \Omega$. Pick a matrix $X \in [0,1]^4$.

$$\begin{array}{cccc}
x_{11} & x_{13} & = z_1 & g(z_1) \\
x_{21} & x_{22} & = z_2 & g(z_2) \\
\| & \| & & \| \\
y_1 & y_2 & & y_0 \\
\hline
(f(y_1) & f(y_2) & & f(y_0))
\end{array}$$

Use g to get the k aggregated opinions: $\big(g(z_1), g(z_2)\big) = y_0 \in [0,1]^2$. Recall that the columns of the matrix X are denoted by (y_1, y_2). Take now any function f and consider the n values that it returns for the n columns $(f(y_1), f(y_2)) = z_0$ and the additional value, $f(y_0)$. By definition, if $g(z_0) = f(y_0)$ then g, f aggregate on X. In other words, if point $(z_0, f(y_0))$ belongs to the graph of function g, i.e., to the set $G = \{(z, g(z)), z \in [0,1]^2\}$ then the two functions aggregate on X. If, on the other hand, g and f do not aggregate on X, then the point $(z_0, f(y_0))$ is not in the graph, i.e., it is in $B = [0,1]^3 \backslash G$. Thus for any matrix X and any given g, a function f can be mapped to a unique point $(f(y_1), f(y_2), f(y_0))$ in $[0,1]^3$, which is either in G or in its complement, B.

The main idea of the proof is that the topological properties of set G are mirrored by those of the set of functions Ω_G that aggregate with g on X.

We impose a very mild restriction on the graph G and thus, on function g : the closure of the graph in $\mathbb{R}^n$ should have an empty interior. Most of the functions one can easily envision satisfy this property: g can be discontinuous and even return multiple values for an isolated set of points. The requirement prevents function g from being "wildly" discontinuous. To construct a function that violates the property, consider the simplest case of g: $\mathbb{R} \to \mathbb{R}$. Partition $[0,1]$ into a countable number of dense (in $[0,1]$) subsets, let g return a different rational number (in $[0,1]$) on each subset, then the closure of the graph of g is $[0,1]^2$, cf. e.g., Drago et al. (2011).

To describe formally the meaning of a small set of functions used as aggregators, we reproduce the definition of a nowhere dense set from Rudin (1991, p. 42).[2]

Definition 2. A subset of a topological space is nowhere dense if its closure in that space has an empty interior.

Finally, we are ready to formulate our main result establishing that for any matrix of opinions and any function g that satisfies the mild requirement discussed above, the set of functions f that aggregates is very small.

Theorem 4.1. *For any matrix $X \in [0,1]^{kn}$ and any Borel function g with the graph $G = \{(z, g(z)), z \in [0,1]^n\}$ whose closure has an empty interior, the set $\Omega_G \subset \Omega$ of Borel functions f such that g, f aggregate on X is nowhere dense in L_∞.*

Proof. Recall that the points $(y_1,\dots,y_n)$ are the columns of matrix X and, since $(z_1,\dots,z_k)$ are the rows of X, point

$$y_0 = \big(g(z_1),\dots,g(z_k)\big)$$

is fully determined by matrix X and function g. So for any fixed g, X, define the family of sets S_t, where $t = (t_1,\dots,t_{n+1}) \in [0,1]^{n+1}$, as a set of functions f that satisfy the $n+1$ conditions:

$$S_t = \{f \in \Omega: f(y_1) = t_1,\dots,f(y_n) = t_n, f(y_0) = t_{n+1}\}$$

Since the functions are Borel, the sets S_t are disjoint. By construction for every function f there is a t such that $f \in S_t$. Thus $(S_t)_{t\in[0,1]^{n+1}}$ is a partition of the set of functions Ω, or a set of equivalence classes. There is a one-to-one map between functions in each equivalence class established by equating the values returned by the functions in all but $n+1$ points: map a function f from equivalence class S_t to a function h in equivalence class S_q if $f(y) = h(y)$ for all $y \notin \{y_1,\dots,y_n,y_0\}$. Thus, the sets S_t are "of the same size".

Next, partition the set $[0,1]^{n+1}$ into two sets, the graph of g, G, and its complement, B. For every $t \in G$, every function f in S_t aggregates with g on X by construction:

$$t \in G \implies g(t_1,\dots,t_n) = t_{n+1}$$
$$f \in S_t \implies f(y_1) = t_1,\dots,f(y_n) = t_n, f(y_0) = t_{n+1}$$

Conversely for any f that aggregates with g on X there is a point q on the graph such that $f \in S_q$. Indeed, by definition of aggregation, $g(f(y_1),\dots,f(y_n)) = f(y_0)$, so let $q_1 = f(y_1),\dots,q_n = f(y_n)$ and $q_{n+1} = f(y_0)$. Clearly then $q \in G$ and $f \in S_q$.

[1] The restriction to Borel functions is natural here: we want to distinguish between functions that differ even at a single point, rather than associating them with the same equivalence class.

[2] A nowhere dense set is an example of a *meagre* (i.e., a very small) set.

Call the set points not on the graph of g, $B = [0, 1]^{n+1} \backslash G$. It follows that f does not aggregate with g on X if and only if $\exists! t \in B: f \in S_t$.

Now we map infinite sequences in $[0, 1]^{n+1}$ to those in Ω. For any sequence $T = \{t_j, j \geq 1: t_j \in [0, 1]^{n+1}\}$ pick the corresponding sequence F which is composed of functions $f_j \in S_{t_j}$, $j \geq 1$ such that for any pair $j, i \geq 1$, $f_j(y) = f_i(y)$ for all $y \notin \{y_1, \ldots, y_n, y_0\}$. Thus the functions in the sequence are equal apart from their values at the $n + 1$ points.

It follows that if the sequence T converges as $j \to \infty$, then F converges as well: values that all $f \in F$ return on $\{y_1, \ldots, y_n, y_0\}$ converge to their corresponding limits, and since the functions are equal elsewhere, they converge in supremum norm.

It also follows that for any converging sequence of t_j in G one can construct the corresponding converging sequence of f that aggregate with g on X, and, in particular, any function $f \in S_t: t \in \bar{G}$, where $\bar{G}$ is the closure of G, is in the closure of the set of functions that aggregate with g on X. It is also easy to see that any function $f \in S_q$ such that $q \notin \bar{G}$ is not in the closure of the set of functions that aggregate. Thus, $\bar{G}$ has an empty interior if and only if the closure of the set of functions f such that g, f aggregate on X has an empty interior. The claim then follows by definition of a nowhere dense set. $\square$

Corollary 4.2. *The set of functions f that aggregate with g consistently (on all X) is substantially smaller, being (an uncountable) intersection of the nowhere dense sets.*

5. Concluding remarks

Our main result shows that the set of aggregators for which the premise- and conclusion-based aggregation give the same answer is scarce. Roughly speaking, starting with an initial pair of aggregators that aggregate consistently and only slightly perturbing one of them (f or g) destroys the consistency. On the other hand, starting with a pair that does not aggregate, and perturbing either function will not help in getting a consistent pair of aggregators.

For a given deliberation rule accepted by all agents and for a given case (represented by their opinions with respect to the basic issues), finding a rule to aggregate these opinions consistently is difficult. This means that a "freely chosen" rule to aggregate individual opinions and a "freely chosen" deliberation rule used to arrive at a conclusion will be subject to the doctrinal paradox, apart from very rare cases. Thus, such choice has to be made carefully, knowing that the set of consistent aggregators is very small.

References

Bovens, L., Rabinowicz, W., 2004. Democracy and argument-tracking truth in complex social decisions. In: van Aaken, C.L.A., Luetge, C. (Eds.), Deliberation and Decision. Ashgate Publishing, Aldershot.

Bovens, L., Rabinowicz, W., 2006. Democratic answers to complex questions-an epistemic perspective. Synthese 150 (1), 131–153.

De Clippel, G., Eliaz, K., 2015. Premise-based versus outcome-based information aggregation. Games Econom. Behav. 89, 34–42.

Dietrich, F., 2006. Judgment aggregation: (im)possibility theorems. J. Econom. Theory 126 (1), 286–298.

Dietrich, F., 2007. A generalised model of judgment aggregation. Soc. Choice Welf. 28 (4), 529–565.

Dietrich, F., List, C., 2007. Arrow's theorem in judgment aggregation. Soc. Choice Welf. 29 (1), 19–33.

Dietrich, F., List, C., 2008. A liberal paradox for judgment aggregation. Soc. Choice Welf. 31 (1), 59–78.

Dietrich, F., List, C., 2010. The aggregation of propositional attitudes: towards a general theory. Oxf. Stud. Epistemol. 3 (215), 34.

Dietrich, F., List, C., 2017a. Probabilistic opinion pooling generalized. Part one: general agendas. Soc. Choice Welf. 48 (4), 747–786.

Dietrich, F., List, C., 2017b. Probabilistic opinion pooling generalized. Part two: the premise-based approach. Soc. Choice Welf. 48 (4), 787–814.

Dietrich, F., Mongin, P., 2010. The premise-based approach to judgment aggregation. J. Econom. Theory 145 (2), 562–582.

Dokow, E., Holzman, R., 2010. Aggregation of binary evaluations. J. Econom. Theory 145 (2), 495–511.

Drago, G., Lamberti, P.D., Toni, P., 2011. A "bouquet" of discontinuous functions for beginners in mathematical analysis. Amer. Math. Monthly 118 (9), 799–811.

Genest, C., Zidek, J.V., 1986. Combining probability distributions: A critique and an annotated bibliography. Statist. Sci. 1 (1), 114–135.

Guilbaud, G.-T., 1952. Les théories de l'intérêt général et le problème de la logique de l'agrégation. Economie appliquée 15, 509–511.

Herzberg, F.S., 2013. Universal algebra for general aggregation theory: Many-valued propositional-attitude aggregators as MV-homomorphisms. J. Logic Comput. 25 (3), 965–977.

Kornhauser, L., Sager, L., 1993. One and the many: Adjudication in collegial courts. Calif. Law Rev. 81, 1–51.

List, C., 2004. A model of path-dependence in decisions over multiple propositions. Amer. Polit. Sci. Rev. 98 (03), 495–513.

List, C., 2005. The probability of inconsistencies in complex collective decisions. Soc. Choice Welf. 24 (1), 3–32.

List, C., Pettit, P., 2002. Aggregating sets of judgements: an impossibility result. Econ. Philos. 18, 89–110.

List, C., Pettit, P., 2004. Aggregating sets of judgments: Two impossibility results compared 1. Synthese 140 (1), 207–235.

McConway, K.J., 1981. Marginalization and linear opinion pools. J. Amer. Statist. Assoc. 76 (374), 410–414.

Nehring, K., 2007. The impossibility of a Paretian rational: A Bayesian perspective. Econom. Lett. 96 (1), 45–50.

Nehring, K., Puppe, C., 2008. Consistent judgement aggregation: the truth-functional case. Soc. Choice Welf. 31 (1), 41–57.

Pauly, M., Van Hees, M., 2006. Logical constraints on judgement aggregation. J. Philos. Logic 35 (6), 569–585.

Pettit, P., 2001. Deliberative democracy and the discursive dilemma. Philos. Issues 11 (1), 268–299.

Rudin, W., 1991. Functional Analysis. McGraw-Hill, Inc., New York.