האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

AN AXIOMATIC APPROACH TO SENSORS' TRUST MEASUREMENTS

By

ALEXANDER KRAVTSOV AND EYAL WINTER

Discussion Paper # 739 (August 2020)

מרכז פדרמן לחקר הרציונליות

THE FEDERMANN CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Edmond J. Safra Campus,
Jerusalem 91904, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

An Axiomatic Approach to Sensors' Trust Measurements

Alexander Kravtsov[1] and Eyal Winter[2]

Abstract

A set of sensors is used to identify which of the users, from a pre-specified set of users, is currently using a device. Each sensor provides a name of a user and a real number representing its level of confidence in the assessment. However, the sensors measure different signals for different traits that are largely unrelated. To be able to implement a policy based on these measurements, one needs to aggregate the information provided by all sensors. We use an axiomatic approach to provide several reasonable trust functions. We show that by providing a few desirable properties we can derive several solutions that are characterized by these properties. Our analysis makes use of an important result by Kolmogorov (1930).

1. Introduction

Security systems for devices need to monitor whether the person using the device for a particular application is authorized to do so. This is typically done by installing a set of sensors in the device that take measurements any time that the device is used and compare these measurements to profiles of a pre-specified set of users. The sensors, however, are very different in terms of the users' traits they are measuring: one might involve voice recognition, a second image recognition, and a third might record body movements. To implement a policy, e.g. to decide whether or not to shut down the device, one needs to aggregate the measurement made by all sensors into a single value for each user that stands for the level of trust that the system assigns to the event that the user is authorized. This is typically done by using a threshold strategy to implement the policy. However, because sensors measure different traits, these measurements may involve different scales, and different units of measurements. Often the exact algorithm designed by the sensor's producer to take these measurements is unrevealed. This makes it hard to impossible to construct a reliable sample space and calculate conditional probabilities that will eventually determine the probabilities of type-one and type-two errors for each policy and identify the optimal policy.

The alternative direction is to use a non-probabilistic measure of similarity between the tested user and each of the users in the pre-specified set of users. Since different sensors use different techniques for aggregating their measurements, there is no straightforward and unique way to

[1] Senior researcher at Toga Network and Huawei, Israel
[2] Silverzweig Professor of Economics at the Hebrew University

do so. This direction begs for an axiomatic approach that will specify a set of desirable properties to be satisfied by the aggregating method. Axiomatic approaches are very prevalent in economic theory and game theory. Perhaps, the most famous among those are von Neumann and Morgenstern (1944) axioms for expected utility theory, Arrow's (1950) Impossibility Theorem, Nash's (1950) bargaining solution, and the Shapley (1953) Value[3] (see also Winter 2002).

At the core of the axiomatic approach lies the idea that instead of coming up with a solution "out of the blue" (where a solution can be a measurement method, a voting system, a fair allocation of resources, and more), it is better to first identify desirable properties (axioms) for a solution and then attempt to characterize all the solutions that satisfy these properties.

In this paper we provide an axiomatic treatment for measuring the trust level of multiple sensors. Our model assumes that each sensor provides two pieces of information: the identity of the user (i.e, which user from among the pre-specified users is likely to be using the device) and a real number that represents the level of confidence in this measurement. This confidence level can be based on a mixture between the sensor's own assessment and an external test. To the set of potential users we add a virtual user called "other" who represents the prediction that the current user is outside the set of pre-specified ones. Our aim is to construct a measurement method that aggregates this data from multiple sensors, with a single number for each user specifying the system's level of trust that the user is the one currently using the device. As mentioned we use an axiomatic approach. We provide three axiomatizations with the first and the third pinning down a unique solution and the second one yielding a set of solutions. Two types of axioms are involved. The first one determines the solution for degenerate cases, i.e., when the system works with a single sensor or when some user is not identified by any of the sensors. But the second, more important, type of axiom deal with the way two data sets (involving different sensors) are merged into one.

We stress here that no set of axioms is immune to potential criticism and we are not claiming that the axiomatic characterizations we provide here are the only reasonable ones. Our main purpose is to highlight several desirable properties and the way they map into solutions to measure security trust and sensor data fusion. We believe that some of these axioms can be modified to develop other solutions to the problem of user identification.

The literature on the theory of sensor fusion and sensor trust is broad and deep, and it is beyond the scope of this paper to review this literature. We point out, however, that much of this literature uses a probabilistic approach. This includes approaches of imprecise probability theories (Shafer's 1976 theory of evidence) or Bayesian approach (Sommer et al 2009). Other approaches are based on the concept of "fuzzy sets"(Durrant-Whyte 2008 and Chen et al 2013) and the theory of subjective logic (Jøsang 2016), which assigns a degree of uncertainty to probabilities. The axiomatic approach in the sensor/information fusion literature is relatively

[3] Arrow (1950), Nash (1950) and Shapley (1953) are all Nobel laureate contributions.

sparse. Dubis et al. (2016) propose an axiomatic treatment for a model of information fusion but their framework is rather abstract. By contrast, our model is simpler and builds on the actual information provided by sensors that are regularly used by many firms in the communication industry. Furthermore, two of our three axiomatizations pin down a unique solution for the problem.

The rest of the paper is organized as follows. Section 2 describes the model framework and provides basic definitions. Section 3 we defines the concept of trust function and presents one such function which is axiomatized in Section 4. Section 5 presents two alternative axiomatizations one of which characterizes a set of solutions that differ in the type of averaging operator that one uses (i.e., arithmetic, harmonic, geometric, etc.) The second axiomatization pins down a unique solution using an axiom that is akin to the way events are assigned probabilities in probability theory. We end with an appendix that provides the proof of an important result by Kolmogorov (1930) that we use for our characterization.

2. Model and Definitions

Let S be a finite set of sensors, and M a finite set of potential users of a device. When a device is in use the sensors have to jointly identify which of the users in M is currently using the device. We add to the set M an additional user "O" ("others") that allows sensors to indicate that the current user is outside the set M, and denote M* = M ⋃ {O}. Sensors provide measurements. A measurement of a sensor i specifies a single user $m_i \in M^*$ (the identified user) and a level of confidence in i's assessment denoted by γ_i . Hence, i's measurement is given by the pair $w_i = (m_i, \gamma_i)$. We assume that the range of confidence levels for all sensors is finite and identical for all sensors[4], i.e., $0 \leq \gamma_i \leq \gamma^*$. We now define the notion of a "data set."

Definition: A sensor data set (SDS) w = $\{w_i\}_{s\epsilon S'} = \{(m_s, \gamma_s)\}_{s\epsilon S'}$ of size n is a sequence of measurements coming from n different sensors in $S' \subset S$ and $|S'| = n$. We denote by W the set of all sensor data sets.

3. Trust Functions

A trust function aggregates the measurements of all sensors to a single number for each user m. It can be interpreted as the level of trust of the entire system in user m being the user who is currently using the device. We stress here that this number does not necessarily reflect a certain probabilistic belief, and hence it does not necessarily lie between 0 and 1. Instead it can

[4] If this is not the case in practice, one can use a normalization; i.e., if γ_i^* is the maximal confidence value for sensor i use [0,1] as the relevant interval for all sensors and set $\gamma_i' = \gamma_i/\gamma_i^*$ for the normalized confidence level.

be viewed as a measure of fitness/similarity between the traits of users, and the overall data that has been collected by the sensors.

Definition: A trust function (TF) is a function F that assigns to each sensor data set $w \epsilon$W a vector in $R_+^{|M^*|}$. We interpret $F_m(w)$ as the level of trust assigned to user m being the true user given the evidence generated by w.

Definition: A trust function F is said to be symmetric if swapping the measurements of two sensors i and j would not affect the value of $F(w)$; i.e., for w and w' with $w_i = w'_j$ and $w'_i = w_j$ for some $i \neq j$ and $w_k = w'_k$ for all $k \neq i, j$ we have $F(w) = F(w')$.

Throughout the paper we will assume symmetry. Obviously in reality one might wish to attach more importance to one sensor than another by giving its measurements a greater weight. However, such asymmetry would typically arise only ex - post, i.e., after the system has been operating sufficiently long to reveal that one sensor's measurements are more accurate than another's.

We shall now present our first proposed trust function:

For each $m \epsilon M$ and a vector w, we denote by $F_m(w)$ the overall trust for m in the following way:

$$F_m(w) = \frac{1}{n} \sum_{j \in S_m} \gamma_j, \qquad (1)$$

where $S_m = \{j | m_j = m\}$.

This most straightforward TF averages the confidence level over all sensors by assigning to all sensors that did not identify the user m a confidence level zero for user m.

Note that (1) can also be written as:

$$F_m(w) = w(m)\mu(w, m), \qquad (2)$$

where $\mu(\mathrm{w}, \mathrm{m}) = \frac{1}{|S_m|} \sum_{j \in S_m} \gamma_j$ and $w(m) = \frac{|S_m|}{n}$ Hence one factor in the product of $F_m(w)$ concerns the propensity of measurements to identify m as the user and the other aggregates the confidence level of these measurements.

Our objective now is to characterize $F_m(w)$ axiomatically and then use the axiomatic approach to propose several other reasonable trust functions.

4. Axiomatization

Our first axiom requires that whether a user is not identified by a sensor or is identified with a confidence level zero, it will have the same effect on the aggregated trust level for that user. Furthermore, if the confidence level for a user m is zero for all sensors then so is the aggregated trust level for m.

Axiom 1 Suppose that w = $\{(m_s, \gamma_s)\}_{s\epsilon S'}$ and w' = $\{(m'_s, \gamma'_s)\}_{s\epsilon S'}$ are TFs that are identical for all but one sensor s' such that $m_{s'} = m$ and $\gamma_{s'} = 0$, and $m'_{s'} \neq m$, then $F_m(w) = F_m(w')$. Furthermore, if for each s such that $m_s = m$, $\gamma_s = 0$, then $F_m(w) = 0$.

Our single sensor axiom requires that the aggregated trust level for identifying a use by a single sensor, is identical to the confidence level for that user.

Axiom 2 Single-sensor condition: If $w = (m, \gamma)$ is an SDS with one sensor, then $F_m(w) = \gamma$.

Axiom 3 concerns the merging of two data sets. It asserts that when two data sets merge the joint trust level has to be a weighted average of the two data sets separately where the weights are proportional to the sizes of the data sets.

Axiom 3 Data expansion: Consider two SDS, $w^1 = \{(m_s, \gamma_s)\}_{s\epsilon S'_1}$ and $w^2 = \{(m_s, \gamma_s)\}_{s\epsilon S'_2}$ of sizes n_1 and n_2 respectively. Then $\frac{n_1}{n_1+n_2} F_m(w^1) + \frac{n_2}{n_1+n_2} F_m(w^2) = F_m(w^1 \cup w^2)$, where $w^1 \cup w^2$ is a SDS of size $n_1 + n_2$ given by $\{(m_s, \gamma_s)\}_{s\epsilon S'_1} \cup \{(m_s, \gamma_s)\}_{s\epsilon S'_2}$.

Proposition 1: There exists a unique trust function F satisfying Axioms 1-3, and it is given by $F_m(w) = w(m)\mu(\mathrm{w}, \mathrm{m})$ for $m\epsilon M$, where n is the size of w.

Proof: We first show that all the axioms are satisfied by the TF F defined in the proposition. For Axiom 1 note that according to (2) a sensor j contributes a positive value to $F_m(w)$ if $j \in S_m$ and $\gamma_j > 0$. Furthermore, from (2) it follows immediately that if $\gamma_s = 0$, for all sensors, then $F_m(w) = 0$. Hence Axiom 1 holds because of the equivalence between (1) and (2). For Axiom 2 note that if w = (m, γ) is a single sensor SDS, then $|S_m| = 1$, $\frac{|S_m|}{n} = 1$, and $\mu(\mathrm{w}, \mathrm{m}) = \gamma$, so we have $F_m(w) = \gamma$ as asserted by Axiom 2.

Finally, we show that the TF defined in the proposition satisfies Axiom 3 by building on the equivalence between (1) and (2). Take two SDS w^1 and w^2 with the set of sensors S^1 and S^2 of size n_1 and n_2 respectively, and let $w^3 = w^1 \cup w^2$, (defined on the set of sensors $S^3 = S^1 \cup S^2$), then

$F_m(w^3) = \frac{1}{n_1+n_2}\left\{\sum_{i\in S_m^1}\gamma_i + \sum_{i\in S_m^2}\gamma_i\right\} = \frac{1}{n_1+n_2}\{n_1\ \frac{|s_m^1|}{n_1}\mu(w^1,\mathrm{m}) + n_2\frac{|S_m^2|}{n_2}\mu(w^2,\mathrm{m})$ as asserted.

We next show uniqueness. We have to show that there is no other TF satisfying Axioms 1-3. We start by showing that Axioms 1 and 2 determine a TF uniquely on SDS of size 1. Suppose by way of contradiction that there are two TFs F and F' satisfying Axioms 1,2 and yet F ≠ F'. If F ≠ F', then there must be some SDS of size 1 w such that $F(w) \neq F'(w)$. Without loss of generality assume that w = (m, γ). By Axiom 2, $F_m(w) = \gamma$. By Axiom 1 for any $m' \neq m$ and $m' \epsilon M$, we have $F_{m'}(w) = 0$. Now because F' satisfies Axioms 1, 2 we would have the same statements for F' contradicting the fact that $F(w) \neq F'(w)$.

We next show that Axiom 1- 3 determine uniquely an SFF on the entire set of SDS: We have to show that if F and F' are both TFs satisfying axioms 1- 3, then we must have $F(w) = F'(w)$ for every SDS w. We show this by induction on the size of the SDS. We use the following induction assumption: For any TF defined on DSD w of size less or equal to r, $F(w) = F'(w)$. We have to show that we must have $F(w) = F'(w)$ for all SDS of size r+1. We have already shown that the statement is true for r=1. Let w^1 be an SDS of size r given by $\{(m_s,\gamma_s)\}_{s\epsilon S_1'}$ with $|S_1'| = r$ and w^2 be an SDS of size 1. We argued earlier in the proof that $F(w^2) = F'(w^2)$. Furthermore, by the induction assumption we also have $F(w^1) = F'(w^1)$. Let now w^3be an SDS of size r+1.
We can write $w^3 = w^1 \cup w^2$, where w^1 is of size r and w^2 is of size 1. Since both F and F' satisfy Axiom 4 we have

$\frac{r}{r+1}F(w^1) + \frac{1}{r+1}F(w^2) = F(w_3)$ and $\frac{r}{r+1}F'(w^1) + \frac{1}{r+1}F'(w^2) = F'(w^3)$, which implies that $F'(w^3) = F(w^3)$. Q.E.D.

5. Trust Functions of Generalized Mean

In this section we shall go beyond the trust function defined and axiomatized in Section 4 by studying more basic aggregation axioms without referring directly to a notion of averaging. As we shall see, this will allow us to provide an axiomatic characterization of a parametrized class of trust functions rather than a single trust function.

We start by strengthening Axiom 2 and require that if all sensors attribute the same level of confidence to a user, the aggregated trust level for that user should coincide with the confidence level.

Axiom 2': Full consensus: Suppose that $w_1 = w_2, \ldots, = w_n = (m,\gamma)$, then $F_m(w) = \gamma$.

Our next axiom is a consistency axiom. It assert that by replacing the measurements of a subset of sensors with their aggregated trust level for user m, the new SDS involving the entire set of sensors will not change.

Axiom 4: Consistency: Suppose that $w = (w_1, w_2, \dots, w_n)$ and $F_m(w_1, w_2, \dots, w_k) = \gamma$ for k ≤ n, then $F_m(w) = (w^*, \dots, w^*, w_{k+1}, w_{k+2}, \dots, w_n)$, where $w^* = (m, \gamma)$.

We can now state our second axiomatization result.

Proposition 2: Suppose that TF F Satisfies Axioms 1, 2', and 4. Then there exists a continuous and increasing real function f, such that

$F_m(w) = f^{-1}(\frac{1}{n}\sum_{j\in S_m} f(\gamma_j))$ for $m \epsilon M$, where n is the size of w.

Proof: We first associate every SDS w = $\{(m_s, \gamma_s)\}_{s\epsilon S'}$ with an n-dimensional real vector of the form $x = (x_1, x_2, \dots, x_n) = (\gamma_1, \gamma_2, \dots \gamma_k, 0, \dots, 0)$, where $\gamma_1, \gamma_2, \dots, \gamma_k$ refer to the trust level of the sensors that identify m and have a positive trust level, and all the coordinates with zero values correspond to either sensors that identify a user $m_j \neq m$ or a sensor that for which $m_j = m$, but with zero confidence level, i.e., $\gamma_j = 0$. Because of the symmetry of F the precise mapping from the coordinates of x to the sensors does not affect the aggregated trust level. By Axiom 1, $x = (x_1, x_2, \dots, x_n)$ contains all the relevant information for $F_m(w)$. We can therefore think of F_m as a function from R_+^n to [0,*], where γ* is the upper bound for the confidence levels. Axiom 2' and Axiom 4 now correspond to properties III and IV in Kolmogorov (1930) (see the Appendix for the proof of the result). Following Kolmogorov (1930) there exists an increasing function f such that

$$F_m(w) = f^{-1}(\frac{1}{n}\sum_{j=1}^{n} f(x_i)), \qquad (3)$$

But $x_i = \gamma_i$ if $i\epsilon S_m$ and $x_i = 0$ if $i\epsilon S \backslash S_m$ where $S_m = \{j | m_j = m\}$.

Hence $F_m(w) = f^{-1}(\frac{1}{n}\sum_{j\in S_m} f(\gamma_j))$, which completes the proof. Q.E.D.

We note here that special cases of (3) include some of the well-known definitions of the mean. In particular,

1. If $f(\gamma_j) = \gamma_j$, then $F_m(w)$ corresponds to the arithmetic mean of the confidence levels i.e., . $F_m(w)$ coincides with the TF as given in Proposition 1.
2. If $f(\gamma_j) = \gamma_j^2$, then $F_m(w)$ corresponds to the quadratic mean, i.e., $F_m(w) = \sqrt{\frac{1}{|S_m|}\sum_{j\in S_m} \gamma_j^2}$

eIf $f(\gamma_j) = \log(\gamma_j)$, then $F_m(w)$ corresponds to the geometric mean, i.e., $F_m(w) = \sqrt[|S_m|]{\prod_{j \in S_m} \gamma_j}$

3. If $f(\gamma_j) = 1/\gamma_j$, then $F_m(w)$ is the harmonic mean of the confidence levels i.e., $F_m(w) = \frac{1}{\frac{1}{|S_m|}\sum_{j \in S_m} 1/\gamma_j}$.

5. A Trust Function of Uncertainty Removal

We now consider replacing Axiom 4 with a different axiom for aggregating sensors' information. This axiom builds on the idea of likelihood. Here, $F_m(w)$ can be interpreted as the degree of removal of evidence to the contrary of m when m is the true user of the device.

Axiom 4' Probability expansion: Consider two SDS, $w^1 = \{(m_s, \gamma_s)\}_{s \epsilon S_1'}$ and $w^2 = \{(m_s, \gamma_s)\}_{s \epsilon S_2'}$ with $S_1' \cap S_2' = \emptyset$, and let $w^3 = w^1 \cup w^2 = \{(m_s, \gamma_s)\}_{s \epsilon S_1'} \cup \{(m_s, \gamma_s)\}_{s \epsilon S_2'}$. Then,

$$F_m(w^3) = F_m(w^1) + F_m(w^2) - F_m(w^1)F_m(w^2),$$

The intuition for Axiom 4' is explained in Figure 1. Consider two sensors 1 and 2 generating the SDS's w^1 and w^2 both of which identify user m. The large rectangular represents the volume of the uncertainty as to whether m is the user. Every sensor with its corresponding SDS contains evidence that removes part of this uncertainty. The vertical small rectangle represents the volume of removed uncertainty by w^1 is, whereas the horizontal rectangle represents the uncertainty removed by w^2. As more and more sensors indicate that the user is m larger regions of the rectangle are covered. If the entire rectangle is covered, then all uncertainty is removed and m is identified with certainty. Since each sensor checks a different aspect of the user's identity we may assume that the sets of removed uncertainty are orthogonal. $F_m(w^1)F_m(w^2)$, now refers to the volume of uncertainty that both sensors remove. Hence, in order to avoid counting this volume twice we have to subtract it from $F_m(w^1) + F_m(w^2)$ to get $F_m(w^3)$.

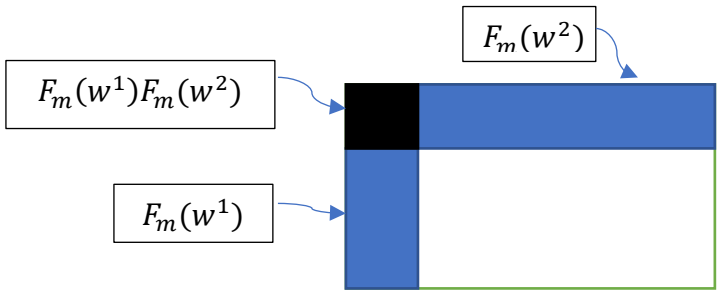

Figure 1

Let $w = \{(m_s, \gamma_s)\}_{s\epsilon S'}$ be an SDS and for each user m we recall that $S_m = \{j | m_j = m\}$.

We now define for each m in M^*

$$G_m(w) = \sum_{k=1}^{|S_m|} (-1)^{k+1} \{ \sum_{\{\forall S \subset S_m, |S|=k\}} \prod_{j\epsilon S} \gamma_j \}$$

$G_m(w)$ is derived by the inclusion-exclusion principle that is used to compute the volume of the union of non-disjoint sets.

Proposition 3 : There exists a unique TF satisfying Axioms 1, 2 and 4', and it is given by $G_m(w)$ defined above.

Proof: We start by noting that Axioms 1 - 2 determines uniquely the trust level for SDSs with a single sensor (see the proof of Proposition 1). The proof of Proposition 3 now follows by induction, namely, by implementing Axiom 4' recursively. Specifically, suppose that there exists another TF F (different from G) satisfying 1,2 and 4'. Let w^1 be an SDS of size r given by $\{(m_s, \gamma_s)\}_{s\epsilon S_1'}$ with $|S_1'| = r$ and w^2 be an SDS of size 1. We argued earlier in the proof that $F(w^2) = G(w^2)$. Furthermore, by induction we also have $F(w^1) = G(w^1)$. Let now w^3 be an SDS of size r+1.

We can write $w^3 = w^1 \cup w^2$, where w^1 is of size r and w^2 is of size 1, and by Axiom 4' we have:

$$F_m(w^3) = F_m(w^1) + F_m(w^2) - F_m(w^1)F_m(w^2) = G_m(w^1) + G_m(w^2) - G_m(w^1)G(w^2)$$
$$= G_m(w^3)$$

Q.E.D.

6. Discussion

1. In this paper we have chose a simple and practical axiomatic approach to deriving sensors' information fusion. We believe that the axioms we introduced here, and especially the aggregation axioms, (i.e., 4 and 4'), can be modified to be used in a more complex framework, where the data contains more than a single value for the confidence level (e.g., both internal and external values,) or when the sensors do not necessarily point to a single user but rather to a set of them. Such an extension might provide new insights and new techniques for aggregating sensors' information.

2. We weren't explicit in the paper regarding the actual meaning of γ_{i}, because different sensors' producers would have different meanings for this variable. However, one aspect of this meaning has to be taken into account. Consider a sensor whose level of confidence ranges between 0 and γ^*, and whose reported level is γ_{i}. There are two ways in which γ^*- γ_{i} can be interpreted. It can represent the degree of uncertainty concerning the assessment "m," or alternatively the degree of certainty that the current user is ***not*** m. Our analysis here was building on the latter interpretation. However, one can quite easily convert it to the first interpretation by distributing γ^*- γ_{i} across all users (including the m). This can be done by invoking a symmetry assumption and dividing equally, or if some prior probabilities are available based on preliminary data, dividing γ^*- γ_{i} proportionally to these priors. An alternative approach to deal with this matter at the practical level is to construct a hybrid model that is a weighted average of the two priors, to test the model with existing data, and then to calibrate it by searching for the weights that best fit the data. This hybrid model approach can also be used to combine different TFs that emerge from our axiomatizations, or even combining one of our TFs with other TFs as proposed in the sensor fusion literature.

3. We finally note that our framework and results are useful also outside the context of sensors' data as the problem described here is akin to many relevant problems in the social sciences. These include the following:

1. Aggregating opinions of experts who are advising on a medical treatment or a technical solution (where m is treatment and γ_j is the extent to which an expert is certain that the treatment is the optimal one.)
2. Aggregating voters' opinions regarding a policy or a candidate (where m is a policy/candidate, and γ_j correspond to the intensity of liking the policy/candidate.)
3. Allocating resources or tasks between individuals (where m is an allocation of resources/tasks and γ_j is the utility an individual is deriving from his proposed allocation).

In all these applications a social planner has to implement a solution based on the preferences reported by individuals. However, the most important difference between an

individual and a sensor, in this context, is that the latter might be gravely inaccurate, but it can never be intentionally dishonest. To be able to apply the techniques developed in this paper to the three applications listed above, one would have to take into account that individuals may intentionally misreport. To avoid such behavior one would need to design a version of our method that is also incentive compatible.

Appendix

Proof of Kolmogorov's Average Characterization Theorem:

Let $M_n(x_1,..,x_k)$ be a symmetric, continuous, and strictly increasing function (in each coordinate) from $R^k \ to \ R$ defined for every $n \geq 1$ and every $x_1,..,x_k, \ a \leq x_i \leq b$. (with a slight abuse of notation we will write $M(x_1,..,x_k)$ omitting the subscript n). We can interpret $x_1,..,x_n$ as an n-dimensional input and $M(x_1,..,x_k)$ as the output aggregating the input to a single value.

We assume the following two properties for M:

I. $M(x,x,x,\dots,x) = x$; i.e., if the input consists of a single value this value must coincide with the aggregated output.

II. $M(x_1,..,x_n) = \bar{x} \quad \Rightarrow \quad M(x_1,..,x_n,y_1,..,y_m) = M(\bar{x},..,\bar{x},y_1,..,y_m)$ i.e., if a subset of the input is aggregated by the value $\bar{x}$, replacing the values in this subset with $\bar{x}$ will yield the same aggregated output.

Theorem (Kolmogorov 1930): A function M satisfies conditions I and II if and only if $M(x_1,..,x_n)$ is of the form (1),

$$M(x_1,x_2,..,x_n) = \phi^{-1}\left(\frac{\phi(x_1)+\phi(x_2)+..+\phi(x_n)}{n}\right), \qquad (1)$$

where ϕ is a continuous and increasing function and ϕ^{-1} is its inverse function.

Proof: The fact that the function in (1) satisfies conditions I and II is straightforward. We will now show that I and II imply that M must be of the form specified in (1). The proof is given for the case where x_i are all rational numbers. The continuous case uses standard approximation arguments together with the continuity of M.

Consider a function M satisfying I and II.

Let $M(mx,ny) = M(x_1,..,x_m,y_1,..y_n)$,

$\boldsymbol{x} = \ x_1 = x_2 =..= x_m, \ \boldsymbol{y} = y_1 = y_2 =..= y_n$ (i.e., $mx \ and \ ny$ are replicas of x and y)

From I and II it follows that $M(pmx,pny) = M\big(pM(mx,ny)\big) = M(mx,ny)$.

If $mn' = nm'$ then

$M(mx, ny) = M(mn'x, nn'y) = M(nm'x, nn'y) = M(m'x, n'y)$. Hence, when x and y, are fixed, the value of $M(mx, ny)$ depends only on the ration between m and n.

We can now define a function $\psi(z)$, for every rational $z = \frac{p}{q}$, $0 \leq z \leq 1$, in the following way:

$$\psi(z) = M(pb, (q-p)a),$$

We will now show that ψ is well defined. Indeed, if $z = \frac{p}{q} = \frac{p'}{q'}$ then $p(q'-p') = p'(q-p)$, and so it holds that $\psi(z)$ has a single value.

We next show that $\psi(z)$, as a function of rational z, is an increasing function.

Let $z' > z$ where z' and z denote fractions with same denominator

$$z' = \frac{p'}{q} \ , \qquad z = \frac{p}{q} \ , \qquad p' > p.$$

By condition I we get: $\psi(z') = M(p'b, (q-p')a) > M(pb, (q-p)a) = \psi(z)$. This inequality follows from the fact that on the LHS the vector in M has fewer "a"s than the vector on the RHS, and the fact that $a < b$.

Since ψ is strictly increasing we can define its inverse function to be $\phi(x)$. ϕ is defined for all values of x (where z is rational). Hence ϕ is an increasing function as well.

Consider again $M(x_1, .., x_n)$ and write $x_i = \psi(z_i)$, where z_i is rational.

We now represent z_i as a ratio of two integers by fixing the value of the denominator across all $1 \leq i \leq$n, i.e., $z_i = \frac{p_i}{q}$. Then $x_i = M(p_i b \ , \ (q-p_i)a)$ and

$M(x_1, x_2, .., x_n) = M(qx_1, qx_2, .., qx_n) = M\big((p_1 + .. + p_n)b, (nq - p_1 - .. - p_n)a\big) =$
$\psi\left(\frac{p_1 + .. + p_n}{nq}\right) = \psi\left(\frac{z_1 + .. + z_n}{n}\right) = \psi\left(\frac{\phi(x_1) + .. + \phi(x_n)}{n}\right)$.

Q.E.D.

References

Arrow, Kenneth J. (1950). "A Difficulty in the Concept of Social Welfare." *Journal of Political Economy.* 58 (4): 328–346.

Chen S, Y. Deng, J. Wu (2013) “Fuzzy Sensor Fusion Based on Evidence Theory and its Application” *Applied Artificial Intelligence* 27 (3) 234-248.

Dubios, Didier, W. Liu, J. Ma, H. Parde (2016) “The basic principles of uncertain information fusion. An organized review of merging rules in different representation frameworks” *Information Fusion,* 32 pp. 12-39

H.F. Durrant-Whyte, T.C. “Henderson, Multisensor data fusion, “in: B. Siciliano, O. Khatib (Eds.), *Handbook of Robotics, Springer*, 2008, pp. 585–610.

Jøsang Audun “Subjective Logic: A Formalism for Reasoning Under Uncertainty,” ISBN *978-3-319-42337-1, Springer Verlag,* 2016.

A.N. Kolmogoroff, (1930) Sur la notion de la moyenne. *Atti Accad. Naz. Lincei Mem. Cl. Sci. Fis. Mat. Natur. Sez*. (6) 12 388–391.

Nash, John (1950). "The Bargaining Problem". *Econometrica*. **18** (2): 155–162.

Neumann, John von and Morgenstern, Oskar, Theory of Games and Economic Behavior. *Princeton, NJ.* 1944.

Shafer Glenn “The Mathematical Theory of Evidence” *Princeton University Press*, 1976.

Shapley, Lloyd S. *"Notes on the n-Person Game -- II: The Value of an n-Person Game"* Santa Monica, Calif*.: RAND Corporation. Princeton University Press*, 1953.

Sommer Klaus-Dieter, Kühn Olaf , Puente León Fernando Siebert, Bernd R.L, . (2009) “A Bayesian approach to information fusion for evaluating the measurement uncertainty” *Robotics and Autonomous Systems* 3, 30, 339-344.

Winter, E., “The Shapley Value,” in *The Handbook of Game Theory, eds. R. J. Aumann and S. Hart, North-Holland* (2002), 2026-2052

H E B R E W U N I V E R S I T Y O F J E R U S A L E M

מרכז פדרמן לחקר הרציונליות

THE FEDERMANN CENTER FOR THE STUDY OF RATIONALITY

List of Recent Discussion Papers

These and earlier papers can be found on our website: www.ratio.huji.ac.il

704. **Sophie bade , yannai a. Gonczarowski**, Gibbard-Satterthwaite Success Stories And Obvious Strategyproofness (October 2016)

705. **Shira Cohen-Zimerman, Ran R. Hassin,** Implicit Motivation Makes the Brain Grow Younger: Improving Executive Functions of Older Adults (January 2017)

706. **Abraham Neyman , Elon Kohlberg**, Cooperative Strategic Games (February 2017)

707. **Ohad Dan, Drorith Hochner- Celnikier, Amy Solnica, Yonatan Loewenstein**, Association Of Catastrophic Neonatal Outcomes With Increased Rate Of Subsequent Cesarean Deliveries (April 2017)

708. **Bezalel Peleg, Shmuel Zamir,** Sequential Aggregation Of Judgments (May 2017)

709. **Johannes Müller-Trede, Shoham Choshen-Hillel, Meir Barneron and Ilan Yaniv,** The Wisdom of Crowds in Matters of Taste (May 2017)

710. **Elon Kohlberg and Abraham Neyman,** Games of Threats (September 2017)

711. **Sergiu Hart,** Repeat Voting: Two-Vote May Lead More People To Vote (October 2017)

712. **Sergiu Hart and Philip J. Reny,** The Better Half of Selling Separately (December 2017)

713. **Shevy Waner, Uzi Motro, Yael Lubin, and Ally R. Harari,** Male mate choice in a sexually cannibalistic widow spider (February 2018)

714. **Maya Bar-Hillel, Tom Noah, and Shane Frederick,** Learning psychology from riddles: The case of stumpers (February 2018)

715. **Bezalel Peleg and Ron Holzman,** Representations of Political Power Structures by Strategically Stable Game Forms: A Survey (February 2018)

716. **Kristoffer Arnsfelt Hansen, Rasmus Ibsen-Jensen, and Abraham Neyman,** The Big Match with a Clock and a Bit of Memory (February 2018)

717. **Bezalel Peleg and Hans Peters,** Self-implementation of social choice correspondences in strong equilibrium (April 2018)

718. **Abraham Neyman,** Additive valuations of streams of payoffs that obey the time-value of money principle: characterization and robust optimization (April 2018)

719. **Bezalel Peleg and Shmuel Zamir,** Judgements aggregation by a sequential majority procedure (June 2018)

720. **Lotem Elber-Dorozko and Yonatan Loewenstein,** Striatial Action-Value Neurons Reconsidered (July 2018)

721. **Orli Bobek, Adiv Gal, David Saltz, and Uzi Motro,** Effect of Nest-Site Microclimatic Conditions on Nesting Success in the Lesser Kestrel (Falco Naumanni) (July 2018)

722. **Barry O'Neill,** Two-Party Agreements as Circular Sets (September 2018)

723. **Adiv Gal, David Saltz, and Uzi Motro,** Effect of Supplemental Feeding on Nesting Success in the Lesser Kestrel (Falco Naumanni) (September 2018)

724. **Maya Bar-Hillel,** The unbearable lightness of self-induced mind corruption (November 2018)

725. **Tomer Siedner,** Optimal pricing by a risk-averse seller (May 2019)

726. **Maya Bar-Hillel and Cass R. Sunstein,** Baffling bathrooms: On navigability and choice architecture (June 2019)

727. **Maya Bar-Hillel and Jacob Lavee,** Lay attitudes toward involuntary organ procurement from death-row prisoners: no, but (June 2019)

728. **Maya Bar-Hillel,** Why didn't I see it earlier? (July 2019)

729. **Maya Bar-Hillel, Tom Noah, and Shane Frederick,** Solving stumpers, CRT and CRAT: Are the abilities related? (October 2019)

730. **Sergiu Hart and Yosef Rinott,** Posterior probabilities: dominance and optimism (November 2019)

731. **Sergiu Hart and Dean P. Forester,** Forecast-hedging and calibration (November 2019)

732. **Maya Bar-Hillel,** The base-rate fallacy in probability judgments (December 2019)

733. **Yigal Attali and Maya Bar-Hillel,** The false allure of fast lures (February 2020)

734. **Uri Zak,** Female Chess Players Do Underperform When Playing Against Men: Commentary on Stafford (2018) (March 2020)

735. **Daniel Kahneman and Maya Bar-Hillel,** Laplace and Cognitive Illusions (June 2020)

736. **Sergiu Hart and Yosef Rinott,** Posterior Probabilities: Nonmonotonicity, Log-Concavity, and Turán's Inequality (July 2020)

737. **Maya Bar-Hillel,** An annotated compendium of stumpers (July 2020).

738. **Alex Gershkov and Eyal Winter,** Exploitative Priority Service (August 2020).

739. **Alexander Kravtsov and Eyal Winter,** An Axiomatic Approach to Sensors' Trust Measurements (August 2020).