Research Statement

Erel Segal-Halevi

2017-01-29

My main goal as a researcher is to develop algorithms that can help people create a better society. This goal leads me to various research topics that combine computer science with the social sciences and the humanities. I have pursued this goal in various ways, both in and out of academy:

1. My Ph.D. research (2013–2016) was about **fair division of land**. I developed algorithms that can help people with different preferences divide land resources among them, such that all participants agree that they received a fair share, without the help of an external assessor.

2. Concurrently (2014–2017), I participated in a research on **maximizing social welfare in markets**. I developed double-auction mechanisms that induce buyers and sellers to reveal their true valuations for items, and perform trade that approximates the maximum gain-from-trade.

3. Concurrently (2015–2017), I participated in a research on **finding equilibria in repeated games**, taking into account the players' risk-attitude. For example, when the players are risk-averse, there are options for cooperation and stable agreements (aka "Nash equilibria") that were not acknowledged so far.

4. In a pre-doctoral position (2011-2013), I worked on **natural language negotiation agents**. I participated in the development of automated agents that can negotiate with humans in English, and can help people improve their negotiation skills.

5. Outside academy, my job in TROS (a start-up company, 2007–2010) was about **trust in social networks**. I participated in developing an algorithm that calculates, for each two members in a social network, how well they know each other through their common friends and friends-of-friends. The algorithm is patented (Rosenthal et al., 2013).

6. As a volunteer administrator in Hebrew Wikisource,[1] I developed scripts that helped people type and present public-domain classic texts. I also participated in workshops on **digital humanities** — using computational text-analysis tools to help researchers in Bible and literature.

7. My M.Sc. thesis (1997–1999) was about **computational linguistics in Hebrew**. I developed the first morphological analyzer and disambiguator for undotted Hebrew texts.

All these topics are interesting and I plan to continue working on them. But to keep this document at a reasonable length, I will focus below on my most recent academic project (point 1). I will present my past work and some of my future plans.

1 Fair Division of Land

Fair division is a research topic that combines economics, political science, mathematics and computer science (Brams and Taylor, 1996; Robertson and Webb, 1998; Brams, 2007; Procaccia, 2015). The general goal is to divide a heterogeneous resource among several partners with different preferences,

[1] https://he.wikisource.org/wiki/User:Erel_Segal

such that each partner receives a share that is "fair" according to his/her own preferences. As a computer scientist, my focus is on the algorithmic aspect of fair division: my goal is to develop procedures that people can use in order to find a fair division, without the need to employ an external "objective" assessor.

During my Ph.D. research I focused on the important special case in which the resource to be divided is *land*. In existing work, land is usually modeled as a *cake*, which is modeled as a 1-dimensional interval. The preferences of the agents are modeled by value-density functions $v_i : Cake \to \mathbb{R}^+$. The value of a piece to an agent is the integral of the value-density: $V_i(X_i) = \int_{x \in X_i} v_i(x)dx$. There are two definitions for fairness: *proportionality* means that each agent gets at least $1/n$ the total cake value in its own eyes: $\forall i : V_i(X_i) \geq V_i(Cake)/n$, and *envy-freeness* means that each agent receives at least as much as any other agent in its own eyes: $\forall i, j : V_i(X_i) \geq V_i(X_j)/n$. There are algorithms for fair division in this model, but unfortunately, they are insufficient for dividing land. This is due to several important considerations that are not important for cakes but are important for land. I present some of these considerations below.

2 Two-dimensional Fair Division

When dividing land, people care about the 2-dimensional geometric shape of their pieces. For example, people prefer a 30-by-30 meters square to a 900-by-1 meters rectangle, even if the integral of their value-density is the same in both cases.

In collaboration with my Ph.D. advisors, Yonatan Aumann and Avinatan Hassidim, we extended the cake model from a 1-dimensional interval to a 2-dimensional subset of the Euclidean plane (such as a square), and added a requirement that the pieces must belong to a certain family S of geometric shapes (such as squares).

In a first work we proved that, in general, it is impossible to guarantee *proportionality*, but it is possible to guarantee a multiplicative approximation to proportionality. For example, when the cake is square and the pieces must be square, it is impossible to guarantee to all agents more than $1/(2n)$ of the cake value, but there is an efficient algorithm that guarantees to each agent at least $1/(4n-4)$ of the cake value. We developed general algorithmic building-blocks that can be combined to create division algorithms for various other shapes of cakes and pieces, such as: fat rectangles, fat polygons and rectilinear polygonal domains (Segal-Halevi et al., 2017). Many interesting geometric settings remain for future work, such as: dividing an unbounded plane, a rectilinear polygon or a convex fat object. Working on these problems may be a good opportunity for me to collaborate with computational geometers.

In a second work, we added to partial proportionality, the additional fairness requirement of *envy-freeness*. As an example result, when the cake is square and the pieces must be square, we developed an algorithm that generates an envy-free division in which, when $n = 2$, each agent gets at least $1/4$ the cake value, and when $n = 3$, each agent gets at least $1/10$ the cake value (Segal-Halevi et al., 2015). Later we extended it to arbitrary n, where each agent gets $1/O(n^2)$ of the cake value. We also developed algorithms for arbitrary fat cakes and fat pieces. In all settings we examined, for $n = 2$, the proportionality attainable in envy-free divisions equals the proportionality attainable in general divisions. This naturally raises the open question of whether this equality always holds? In other words, does envy-freeness has a "cost" in terms of proportionality? This question may have interesting economic and psychological implications.

The work on 2-dimensional envy-free division created an interesting offshoot related to 1-dimensional cake-cutting. In two dimensions, it is clear that some cake must remain undivided (for example, when dividing a square cake to two square pieces, some cake will inevitably remain undivided). But in one dimension, it is common to assume that all cake must be divided. Moreover, there is a famous impossibility result showing that an envy-free division with connected pieces (intervals) cannot be found in finite time (Stromquist, 2008), but this result relies on the assumption that the entire cake must be divided. We conjectured that if, instead of requiring that the entire cake is divided, we only require proportionality, we can overcome this impossibility result. Indeed, we proved that for 3 agents, we can find an envy-free and proportional division of an interval to connected intervals in finite time. For four or more agents, we developed finite-time algorithms for envy-free division but they are only partially-proportional (Segal-Halevi et al., 2016). Later, Aziz and Mackenzie (2016) improved on our results and presented finite-time algorithms for envy-free division with connected pieces in which each agent receives at least $1/(3n)$ of the total value. It is still an open question whether there exists a finite-time algorithm for envy-free and proportional division for 4 or more agents.

3 Re-division and Land Reform

While a cake is usually divided once and for all, land often has to be re-divided. This can happen because new land-resources are found or because the population changes or because the existing division is considered unfair. I studied various aspects of this re-division process.

First, from an economic perspective, it is important to have division rules that are both fair and *monotonic*. Monotonicity means that, when there is more land to divide, the value of all agents weakly increases; no agent loses from growth. Together with Balzs Sziklai, an economist from Budapest, we proved that the division rule that maximizes the *product* of utilities (also called the *Nash rule*, after the Nash bargaining solution (Nash, 1950)) is proportional, envy-free, Pareto-efficient and monotonic (Sziklai and Segal-Halevi, 2015). We also proved that this rule is essentially the only rule in a large class of natural rules that has these properties. However, the Nash rule does not guarantee that the pieces are connected, and indeed, when the pieces must be connected, we have a monotonic algorithm only for two people. An interesting question for future work is whether there exists a fair and monotonic rule for three or more agents with connected pieces.

Second, from a political perspective, the existing division might be considered unfair, and the government may want to re-divide it in a fair way. This process is called a *land reform*, and it happened many times throughout history and in many countries around the globe (Powelson, 1988). A modern land reform is currently going on in Scotland (Bryden and Geisler, 2007; Hoffman, 2013). The main difficulty in land reform is the need o balance the fairness requirements with the ownership rights of the current landlords. To define this balance formally, I defined quantitative measures of fairness and ownership: *f-fairness*, for some $f \in [0, 1]$, means that each agent receives at least f/n the total cake value; *w-ownership*, for some $w \in [0, 1]$, means that each agent receives at leas w of his/her value in the initial division. I developed an algorithm for finding a division that simultaneously satisfies f-fairness and w-ownership whenever $f + w \leq 1$, and proved that when $f + w > 1$, there might be instances in which such a division is not possible. This algorithm can be seen as a procedure for land reform: once the constants f, w are determined by the government (based on moral or political considerations), the corresponding division can be calculated efficiently. I also proved that such a division is not possible when the pieces must satisfy a geometric constraint such as connectivity (in one dimension), rectangularity or convexity (in two dimensions). For these cases, I presented algorithms that achieve a weaker notion of ownership called *democratic ownership*. I proved that democratic ownership can be attained together with 1/3-fairness (with interval pieces), 1/4-fairness (with rectangular pieces) and 1/5-fairness (with convex pieces). A main question that remains for future work is whether these fractions can be improved using more sophisticated algorithms (Segal-Halevi, 2016).

The work on two-dimensional re-division created an offshoot work in computational geometry that is interesting in its own right. Suppose we have a rectilinear polygonal land-estate in which some rectangular land-plots are already allocated, but some parts of the estate are empty. We would like to create a complete partition of the cake to rectangular pieces such that (1) each existing land-plot is completely contained in a new land-plot, and (2) the number of "blanks" (new land-plots) is minimized. In collaboration with Arseniy Akopyan, a mathematician from Austria, we developed tight bounds on the optimal number of blanks in this and other similar two-dimensional settings (Akopyan and Segal-Halevi, 2016). A main question for future work is what happens when the cake and the pieces have three or more dimensions.

4 Future topic: experiments in fair division

Besides the open problems from past work, I would like to study fair division of land from an empirical perspective. There are surprisingly few experiments that test the performance of fair division algorithms in practice, most of which are about indivisible items (Dupuis-Roy and Gosselin, 2009; Daniel and Parco, 2005; Gal et al., 2016) only one is about cake-cutting (Walsh, 2011). If fair division algorithms are to be used in practice, they should be tested extensively on real data. I plan to do two kinds of experiments.

1. Simulations with land-value data. In a preliminary experiment, I took a map with detailed information on market-value of various land-spots in New Zealand. I added to each value some uniformly-distributed random noise, in order to simulate the fact that different people that have different valuations that are correlated around the market prices. I converted the map to a 1-dimensional interval by converting each column to a single piece whose value is the sum of the column values. I compared two algorithms: (a) the division algorithm commonly used by an assessor, who gives each person a plot with the same market-value, ignoring the personal valuations; (b) the Even-Paz algorithm (Even and Paz, 1984), which gives each person a plot with a subjective value of at least $1/n$. I compared these algorithms in terms of utilitarian social welfare (average utility per agent), egalitarian social welfare (smallest utility per agent), and maximal envy per agent. The Even-Paz algorithm was significantly better than the assessor algorithm in all these parameters, and the difference increased linearly with the amplitude of the noise and with $\log(n)$. There are many more issues that I would like to check, for example:

- Verify the results on value-maps in different countries.
- Compare different division algorithms, such as the Last Diminisher (Steinhaus, 1948).
- Check the effect of post-processing the algorithm results, such as with the Top Trading Cycles algorithm (Shapley and Scarf, 1974). How much social welfare can be gained?
- It is known that cake-cutting algorithms are not truthful, so agents can gain by misreporting their preferences. But how much exactly they can gain? Specifically, suppose $n-1$ agents play truthfully, and the n-th agent knows their entire valuation function and reports the valuation function that is best for its own interests — how much will this agent gain relative to the truthful outcome?
- What happens when the Even-Paz algorithm is executed on the two-dimensional map (rather than the 1-dimensional interval), where in each step we cut the cake in its longer side? How thin are the final pieces?
- How do our specialized two-dimensional algorithms (Segal-Halevi et al., 2017) fare in comparison with the Even-Paz and the assessor algorithms?

The experiments will probably lead to new interesting questions.

2. Experiments with human subjects. In some of the talks I gave on fair cake-cutting, I asked volunteers from the audience to come and 'divide' among them some land-estate (e.g. a satellite picture of the university). I noticed that people find it hard to understand what they should do. E.g, they do not know how to cut a piece "exactly in the middle" (as they should do by the Even-Paz protocol). These very preliminary experiments indicate that we can learn a lot from experimenting with human subjects. For example:

- What algorithm is easiest for the participants to understand?
- What algorithm gives the participants the highest subjective feeling of "fairness"?
- To what extend do people use strategic considerations in dividing the cake?

Experimenting with human subjects raises several dilemmas. First, how should we encode the participants' preferences: should we let them play by their own subjective preferences regarding the divided resource, as done e.g. by Gal et al. (2016)? Or should we tell them what their valuations are and incentivize them by paying them money according to their value-points, as done e.g. by Herreiner and Puppe (2009)? Second, should aim for a large international audience through a website, like in the experiments of (Eliaz and Rubinstein, 2011)? Or should we experiment with a small, controlled group of students in the laboratory, as done e.g. by Schneider and Krämer (2004)? Each method has its pros and cons, and I will be happy to work with partners experienced in this kind of experimental work.

References

Akopyan, A. and Segal-Halevi, E. (2016). Cutting a Cake without Harming the Toppings. arXiv preprint 1604.00960.

Aziz, H. and Mackenzie, S. (2016). A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *FOCS 2016*.

Brams, S. J. (2007). *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press, first edition edition.

Brams, S. J. and Taylor, A. D. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

Bryden, J. and Geisler, C. (2007). Community-based land reform: Lessons from Scotland. *Land Use Policy*, 24(1):24–34.

Daniel, T. and Parco, J. (2005). Fair, Efficient and Envy-Free Bargaining: An Experimental Test of the Brams-Taylor Adjusted Winner Mechanism. *Group Decision and Negotiation*, 14(3):241–264.

Dupuis-Roy, N. and Gosselin, F. (2009). An Empirical Evaluation of Fair-Division Algorithms. In *CogSci*.

Eliaz, K. and Rubinstein, A. (2011). Edgar Allan Poe's riddle: Framing effects in repeated matching pennies games. *Games and Economic Behavior*, 71(1):88–99.

Even, S. and Paz, A. (1984). A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296.

Gal, Y. K., Mash, M., Procaccia, A. D., and Zick, Y. (2016). Which Is the Fairest (Rent Division) of Them All? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 67–84, New York, NY, USA. ACM.

Herreiner, D. and Puppe, C. (2009). Envy Freeness in Experimental Fair Division Problems. 67(1):65–100.

Hoffman, M. (2013). Why community ownership? Understanding land reform in Scotland. *Land Use Policy*, 31:289–297.

Nash, J. F. (1950). The Bargaining Problem. *Econometrica*, 18(2).

Powelson, J. P. (1988). *The Story of Land: A World History of Land Tenure and Agrarian Reform*. Lincoln Institute of Land Policy.

Procaccia, A. D. (2015). Cake Cutting Algorithms. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, chapter 13, pages 261–283. Cambridge University Press.

Robertson, J. M. and Webb, W. A. (1998). *Cake-Cutting Algorithms: Be Fair if You Can*. A K Peters/CRC Press, first edition.

Rosenthal, T., Segal-haLevi, E., and Segal, E. (2013). Us patent 8364605: Calculating connectivity, social proximity and trust level between web user.

Schneider, G. and Krämer, U. S. (2004). The Limitations of Fair Division. *Journal of Conflict Resolution*, 48(4):506–524.

Segal-Halevi, E. (2016). How to Re-Divide a Cake Fairly. arXiv preprint 1603.00286.

Segal-Halevi, E., Hassidim, A., and Aumann, Y. (2015). Envy-Free Cake-Cutting in Two Dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 1021–1028. arXiv preprint 1609.03938.

Segal-Halevi, E., Hassidim, A., and Aumann, Y. (2016). Waste Makes Haste: Bounded Time Algorithms for Envy-Free Cake Cutting with Free Disposal. *ACM Trans. Algorithms*, 13(1):12:1–12:32.

Segal-Halevi, E., Nitzan, S., Hassidim, A., and Aumann, Y. (2017). Fair and Square: Cake-Cutting in Two Dimensions. *Minor revision for Journal of Mathematical Economics.* arXiv preprint 1510.03170.

Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37.

Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1):101–104.

Stromquist, W. (2008). Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics.*

Sziklai, B. and Segal-Halevi, E. (2015). Resource-monotonicity and Population-monotonicity in Cake-cutting.

Walsh, T. (2011). Online Cake Cutting. *Algorithmic Decision Theory*, 6992:292–305.