



Rental Harmony: Sperner's Lemma in Fair Division

Francis Edward Su

The American Mathematical Monthly, Vol. 106, No. 10. (Dec., 1999), pp. 930-942.

Stable URL:

<http://links.jstor.org/sici?sici=0002-9890%28199912%29106%3A10%3C930%3ARHSLIF%3E2.0.CO%3B2-L>

The American Mathematical Monthly is currently published by Mathematical Association of America.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/maa.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

Rental Harmony: Sperner's Lemma in Fair Division

Francis Edward Su

My friend's dilemma was a practical question that mathematics could answer, both elegantly and constructively. He and his housemates were moving to a house with rooms of various sizes and features, and were having trouble deciding who should get which room and for what part of the total rent. He asked, "Do you think there's always a way to partition the rent so that each person prefers a different room?"

As we shall see, with mild assumptions, the answer is yes. This rent-partitioning problem is really a kind of *fair-division* question. It can be viewed as a generalization of the age-old *cake-cutting* problem, in which one seeks to divide a cake fairly among several people, and the *chore-division* problem, posed by Martin Gardner in [6, p. 124], in which one seeks to fairly divide an undesirable entity, such as a list of chores. Lately, there has been much interest in fair division (see the recent books [3] and [11]), and each of the related problems has been treated before (see [1], [4], [10]).

We wish to explain a powerful approach to fair-division questions that unifies these problems and provides new methods for achieving approximate *envy-free* divisions, in which each person feels she received the "best" share. This approach was carried out by Forest Simmons [12] for cake-cutting and depends on a simple combinatorial result known as Sperner's lemma. We show that the Sperner's lemma approach can be adapted to treat chore division and rent-partitioning as well, and it generalizes easily to any number of players.

From a pedagogical perspective, this approach provides a nice, elementary demonstration of how ideas from many pure disciplines—combinatorics, topology, and analysis—can combine to address a real-world problem. Better yet, the proofs can be converted into constructive fair-division procedures.

1. SPERNER'S LEMMA FOR TRIANGLES. Our fair division approach is based on a simple combinatorial lemma, due to Sperner [13] in 1928. However, do not be fooled—this little lemma is as powerful as it is simple. It can, for instance, be used to give a short, elementary proof of the Brouwer fixed point theorem [7].

As motivation, we examine a special case of Sperner's lemma. Consider a triangle T triangulated into many smaller triangles, called *elementary* triangles, whose vertices are labelled by 1's, 2's, and 3's, as in Figure 1.

The labelling we have chosen obeys two conditions: (1) all of the main vertices of T have different labels, and (2) the label of a vertex along any edge of T matches the label of one of the main vertices spanning that edge; labels in the interior of T are arbitrary. Any labelled triangulation of T satisfying these conditions is called a *Sperner labelling*. The claim:

Sperner's Lemma for Triangles. *Any Sperner-labelled triangulation of T must contain an odd number of elementary triangles possessing all labels. In particular, there is at least one.*




Figure 1. A Sperner labelling, with (1,2,3)-triangles marked.

In Figure 1, we have marked all elementary 123-triangles; their parity is indeed odd. An analogous statement holds in any dimension, which we develop presently.

2. THE n -DIMENSIONAL SPERNER'S LEMMA. We need the concept of an n -simplex: a 0-simplex is a point, a 1-simplex is a line segment, a 2-simplex is a triangle, a 3-simplex is a tetrahedron, etc. In general, an n -simplex may be regarded as an n -dimensional "tetrahedron"—the convex hull of $n + 1$ affinely independent points in \mathbf{R}^m , for $m \geq n$. These points form the vertices of the simplex. A k -face of an n -simplex is the k -simplex formed by the span of any subset of $k + 1$ vertices.

A *triangulation* of an n -simplex S is a collection of (distinct) smaller n -simplices whose union is S , with the property that any two of them intersect in a face common to both, or not at all. The smaller n -simplices are called *elementary simplices*, and their vertices are called *vertices of the triangulation*.

Given an n -simplex S , any face spanned by n of the $n + 1$ vertices of S is called a *facet*. As examples, the facets of a line segment are its endpoints, the facets of a triangle are its sides, and the facets of a tetrahedron are its triangular faces.

Now number the facets of S by $1, 2, \dots, n + 1$. Given a triangulation of S , consider a labelling that obeys the following rule: each vertex is labelled by one of the facet numbers in such a way that on the boundary of S , none of the vertices on facet j is labelled j . The interior vertices can be labeled by any of the facet numbers. Such a labelling is called a *Sperner labelling of an n -simplex*; it generalizes the definition we encountered earlier for $n = 2$. For other low dimensions, Figures 2 and 3 show examples of a Sperner-labelled 1-simplex and 3-simplex.




Figure 2. A triangulated line, with Sperner labelling.




Figure 4. House, rooms, and doors indicated by dotted lines.

carries the first n of the $n + 1$ labels. In our running example, doors are $(1, 2)$ -edges that may be in the interior or on the boundary; see Figure 4. For the case $n = 3$, doors are any room facets labelled $(1, 2, 3)$.

We claim that the number of doors on the boundary of S is odd. Why? The only facet that can contain doors is the $(n + 1)$ -st because of the Sperner labelling. But that facet of S is Sperner-labelled using the labels $1, \dots, n$, hence by the inductive hypothesis there must be odd number of fully labelled $(n - 1)$ -simplices on that facet. These are boundary doors when considered in S .

The boundary doors can be used to locate fully labelled rooms by what we fondly call a “trap-door” argument. The key observation is that every room can have at most 2 doors, and it has exactly 1 door if and only if the room is fully labelled in S . This is true because any room with at least one door has either no repeated labels (it is fully labelled), or it has one repeated label that appears twice. These give rise to 2 distinct doors, one for each repeated label. As examples, verify that elementary triangles in Figure 4 have either two, one, or no $(1, 2)$ -edges. For $n = 3$, verify that a tetrahedron with labels $\{1, 2, 3, 3\}$ has two doors.

So, start at any door on the boundary (located by the inductive step), and “walk” through the door into the adjoining room. Either this room is fully labelled or it has one other door—a “trap-door” that we can walk through. Repeat this procedure, walking through doors whenever possible. Notice that this path cannot double back on itself (because each room has at most two doors), so no room is ever visited twice. Moreover the number of rooms is finite and so the procedure must end, either by walking into a fully labelled room or by walking back through to a boundary door of S ; see Figure 5.

Since the number of boundary doors of S is odd, and trap-door paths pair up only an even number of them, the number of boundary doors left over that lead to fully labelled rooms must be odd. Moreover, any fully labelled rooms not reachable by paths from the boundary must come in pairs, matched up by their own trap-door paths, as in Figure 5. Hence the total number of fully labelled rooms in S is odd. ■




Figure 5. Walking through doors.

This proof yields a constructive method for finding such rooms in the following way. Trap-door paths in successive dimensions can be linked up at their endpoints, because a fully labelled room in an i -dimensional face is just a boundary door in an $(i + 1)$ -face. This creates “super-paths” with endpoints in the bottom and top dimensions, i.e., either $(1, 2)$ -edges on a 1-face of S , or n -dimensional fully labelled rooms in the interior. The constructive procedure begins by moving along the 1-face of S spanned by labels 1 and 2, following any super-path that is encountered. Because the number of $(1, 2)$ -edges is odd, and super-paths can pair up only an even number of them, we see that at least one super-path can be followed to yield a fully labelled room.

The trap-door argument to prove Sperner’s lemma constructively dates back to Cohen [5] and Kuhn [8]. A quick non-constructive proof would note the equality between the number of doors in each room, summed over all rooms, and the number of times each door is counted, summed over all doors. Modulo two, the first sum captures the parity of the number of fully labelled rooms, and the second sum captures the parity of the number of boundary doors, which by the inductive hypothesis is odd.

3. SIMMONS’ APPROACH TO CAKE-CUTTING. Now imagine a rectangular cake to be divided among n people, who may have differing notions of what is valuable on a cake. We use $n - 1$ knives to cut along planes parallel to the left edge of the cake, as in Figure 6.




Figure 6. A cut-set of a cake.

The set of cuts is fully defined by the relative sizes of the pieces. Assume that the total size of the cake is 1 and denote the physical size of the i -th piece by x_i ; this is an absolute measure, unrelated to player preferences. Thus $x_1 + x_2 + \dots + x_n = 1$ and each $x_i \geq 0$. The space S of possible partitions naturally forms a standard $(n - 1)$ -simplex in \mathbf{R}^n . Each point in S corresponds to a partition of the cake by a set of cuts, which we shall call a *cut-set*.

Given a cut-set, we say that a player *prefers* a given piece if the player does not think any other piece is better. We assume that this preference depends on the player and the entire cut-set, but not on choices made by the other players. Note that, given a cut-set, a player always prefers at least one piece, and may (in case of ties) prefer more than one piece by our definition.

We make the following two assumptions:

- (1) *The players are hungry.* That is, players prefer any piece with mass to an empty piece.
- (2) *Preference sets are closed.* This means that any piece that is preferred for a convergent sequence of cut-sets is preferred at the limiting cut-set. Note that this condition rules out the existence of single points of cake with positive desirability.

Theorem. *For hungry players with closed preference sets, there exists an envy-free cake division, i.e., a cut-set for which each person prefers a different piece.*

We first investigate what happens for $n = 3$ people. Suppose the players are named Alice, Betty, and Charlie. They are to divide a cake of total size 1, using 2 knives. Denote the physical size of the pieces by x_1, x_2, x_3 . Since $x_1 + x_2 + x_3 = 1$ and all $x_i \geq 0$, the solution space S is a plane intersected with the first octant. This is just a triangle.

Now triangulate S and assign “ownership” to each of the vertices as in Figure 7, where A stands for Alice, B for Betty, and C for Charlie. We have purposely




Figure 7. Labelling by ownership.

assigned ownership so that each elementary triangle is an ABC triangle. Observe that a similar triangulation of finer mesh can also be labelled in this way.

We obtain a *new* auxiliary labelling of the triangulation by 1's, 2's, and 3's by doing the following: since each point in the triangle corresponds to a set of cuts of cake, go to each vertex, and ask the owner of that vertex, "Which piece would you choose if the cake were cut with this cut-set?" Label that vertex by the number of the piece that is desired.

We claim that this new labelling is a Sperner labelling! Why?

At the vertex $(1,0,0)$ of S we see that one of the pieces contains the entire cake, and the other pieces are empty. By the hungry assumption, the owner of $(1,0,0)$ always chooses piece 1 no matter who the owner is. Similarly $(0,1,0)$ is labelled 2 and $(0,0,1)$ is labelled 3. Next, observe that the sides of the triangle correspond to cuts in which one piece is devoid of any cake. Because no one would ever choose this empty piece, each side of S is missing one label corresponding to the piece that is empty. Hence the Sperner labelling condition is satisfied.

By Sperner's lemma, there must be a $(1, 2, 3)$ -elementary simplex in the triangulation. Since every such simplex arose from an ABC triangle, this means that we have found 3 very similar cut-sets in which different people choose different pieces of cake.

To show the existence of a single cut-set that would satisfy everyone with different pieces, carry out this procedure for a sequence of finer and finer triangulations, each time yielding smaller and smaller $(1, 2, 3)$ -triangles. By compactness of the triangle and decreasing size of the triangles, there must be a convergent subsequence of triangles converging to a single point. Such a point corresponds to a cut-set in which the players are satisfied with different pieces. Why?

Since each $(1, 2, 3)$ -triangle in the convergent subsequence arises from an ABC triangle, consider the choices that the players made in each. With only finitely many ways for players to choose pieces, there must be an infinite subsequence in which the choices of A , B , and C are all constant. Closed preference sets guarantee that at the limit point of this subsequence of triangles, the players are satisfied with distinct choices.

4. THE n -PLAYER CASE. The preceding proof generalizes easily for n players. The only issue that must be addressed is the choice of triangulation for S when $n > 3$. We need a triangulation in which each elementary simplex can be fully labelled by the names of the players. The triangulation we proposed for $n = 3$ does not generalize easily. However, one that works for arbitrary dimensions is a triangulation by *barycentric subdivision*. Loosely speaking, this procedure takes each elementary simplex in a triangulation and subdivides it by marking the barycenters of the faces in each dimension and connecting them to form a new triangulation. A rigorous description of this procedure may be found in [15]. Observe that the mesh of this triangulation can be made arbitrarily small by iterating this procedure; see Figure 8.

Suppose we have iterated barycentric subdivision m times. The desired labelling can be achieved by allowing all vertices that remain from the $(m - 1)$ -th iteration to be labelled A . Any new vertices introduced in the m -th barycentric subdivision are barycenters of simplices of the $(m - 1)$ -th subdivision. To each class of vertices that are barycenters of faces of the same dimension, assign a distinct owner from the persons remaining. There are $n - 1$ such classes. One may verify that this fully




Figure 8. Barycentric subdivision in a 2-simplex, iterated twice.

labels each of the elementary simplices by owners, because each edge connects vertices of different classes.

Now the proof continues almost exactly as in the case $n = 3$: since each point in S corresponds to a cut-set, we construct a new labelling of the triangulation by asking the owner of each vertex, “Which piece would you choose if the cake were cut with these cuts?” The new auxiliary labelling is a Sperner labelling and yields nearby cut-sets that satisfy each person differently. Because this may be done with arbitrarily fine triangulations, by taking subsequences, one may find sequences of cut-sets all converging to one set of cuts in which each person chooses a different piece.

5. A CONSTRUCTIVE APPROXIMATE ALGORITHM. Notice that the preceding proof yields a constructive ϵ -approximate algorithm for cake-cutting—namely, for any prespecified ϵ (such as at the level of crumbs), one may find a set of cuts in which each person receives a piece he considers to be the best up to ϵ -tolerance in the size of the pieces. Simply start the procedure with triangulation mesh size less than ϵ , and then the “trap-door” argument gives a constructive method for finding a fully owner-labelled elementary simplex. Choosing any vertex of this simplex yields a cut-set representing the desired ϵ -approximate solution.

Such an algorithm could be implemented on a computer, which could keep track of what cuts to suggest tentatively and which player to ask, by simply following trap-doors through the simplex of cut-sets. Note that players do not have to state their preferences on every vertex in the triangulation, but only on vertices near a trap-door path, i.e., the complete auxiliary labelling may not need to be determined. So while this algorithm terminates in a number of steps bounded by the number of simplices of the triangulation, it can terminate much sooner.

We emphasize that this notion of ϵ -approximation is based on the physical size of the pieces, not on any quantitative measure of player preferences. However, if one assumes the players’ measures are continuous over the simplex, then by compactness of the simplex and the finite number of players, for any $\epsilon > 0$ there exists a $\delta > 0$ such that pieces of physical size less than δ are believed by each of the players to be size less than ϵ .

6. CHORES AND RENT-PARTITIONING. Now we show how Simmons’ cake-cutting method can be adapted to address other fair-division problems, such as chore division and rent-partitioning.

Finding schemes for envy-free chore division has historically been a more complicated problem than cake-cutting. Most envy-free procedures for cake-cutting do not carry over to chore division without significant modifications. Oskui

[9] solved the case for 3 people; following modifications proposed by Brams and Taylor in [2, pp. 37–39] and [3, pp. 153–55], Peterson and Su [10] gave an explicit chore division scheme for an arbitrary number of players. We now give a simpler ϵ -approximate algorithm for chore division, which falls out nicely as a special case of the *rent-partitioning* problem.

In this problem, n housemates have decided to rent an n -bedroom house for some fixed rent. Each housemate may have different preferences—one may prefer a large room, another may prefer a room with a view, etc. Is there a method for fairly dividing the rent among the rooms? We prove the following:

Rental Harmony Theorem. *Suppose n housemates in an n -bedroom house seek to decide who gets which room and for what part of the total rent. Also, suppose that the following conditions hold:*

- (1) **(Good House)** *In any partition of the rent, each person finds some room acceptable.*
- (2) **(Miserly Tenants)** *Each person always prefers a free room (one that costs no rent) to a non-free room.*
- (3) **(Closed Preference Sets)** *A person who prefers a room for a convergent sequence of prices prefers that room at the limiting price.*

Then there exists a partition of the rent so that each person prefers a different room.

Condition (1) ensures that the problem is well-posed—one cannot talk about *preferences* if some person finds no room *acceptable*, which might happen, for instance, if the rent is too high for all rooms or the rooms are in poor condition.

The miserly condition (2) can be relaxed a bit, as we show in Section 8. The condition also rules out “free closets,” i.e., rooms in which no one would live, even if free.

Condition (3) merely says that in the space of all pricing schemes, preference sets are closed in the topological sense. Note that preference sets may overlap—if in some pricing scheme a person equally prefers two rooms, that person can be assigned to either room.

The rent-partitioning problem may be viewed as a generalization of the cake-cutting problem, in which one seeks to divide *goods* fairly, and the chore division problem, in which one seeks to divide *bads* fairly. However, since the rooms (the goods) are indivisible, known cake-cutting solutions cannot be applied to this problem. And since the rental payments (the bads) are attached to specific rooms, they cannot be divided into more than n pieces and reassembled, which rules out the use of known envy-free chore-division methods such as the discrete method proposed in [3, pp. 154–55] and the procedures proposed in [10]. The two other moving knife schemes proposed for chore division in [3, pp. 153–54] guarantee each player at most $1/n$ of the chores, but are not envy-free.

Alkan, Demange, and Gale [1, pp. 1031–32] have addressed this generalization directly and offer a solution to rent-partitioning via constrained optimization. They implicitly assume conditions equivalent to our conditions (1) and (3), and use a condition weaker than condition (2), but not quite as weak as the condition (2') that we give in Section 8.

We now show how a Sperner’s lemma approach can address the rent-partitioning problem.

7. RENTAL HARMONY: CAKE-CUTTING WITH A TWIST. Our proof of the Rental Harmony Theorem follows Simmons' proof for cake-cutting, but with a twist, so we sketch it.

Suppose there are n housemates, and n rooms to assign, numbered $1, \dots, n$. Let x_i denote the price of the i -th room, and suppose that the total rent is 1. Then $x_1 + x_2 + \dots + x_n = 1$ and $x_i \geq 0$. From this we see that the set of all pricing schemes S forms an $(n - 1)$ -simplex in \mathbf{R}^n .

Now triangulate this simplex by barycentric subdivision of small mesh size. Label it with a fully labelled vertex labelling by the names of the housemates (the same scheme as suggested for cake-cutting). The name at each vertex will be considered the "owner" of that vertex; recall that each vertex corresponds to some pricing scheme for the rooms.

Construct a new labelling from the old by asking the owner at each vertex in the triangulation: "If the rent were to be divided according to this pricing scheme, which room would you choose?" Condition (1) ensures that some answer can be given. Label the vertex by the number of the room that is answered. Let ties in preference be broken arbitrarily.

Here's the twist: the new labelling that results is quite different from the one that arose in cake-cutting. It is not a Sperner-labelling. However, because of the miserly condition (2), it has the property that along each $(n - k)$ -dimensional face, k rooms are free and thus owners along that face prefer one of those k rooms. Figure 9 shows what such a labelling looks like for $n = 3$.

Is there a Sperner-like combinatorial lemma that shows the existence of a fully labelled elementary simplex in this triangulation?

If so, one could proceed as in cake-cutting, by taking finer and finer triangulations to get a sequence of fully labelled elementary simplices converging to a point,




Figure 9. The dual labelling arising from rent-partitioning.

which by condition (3) yields a pricing scheme in which all housemates prefer different rooms. So, all that remains is to establish the Sperner-like combinatorial lemma with a constructive proof.

There are two ways one may proceed. The reader may enjoy proving a Sperner-like lemma for this labelling by using a trap-door argument. The interesting thing that one discovers about this labelling is that on each facet, there is only *one* fully labelled simplex that can be followed into the interior, so that the trap-door procedure succeeds without returning to the boundary again.

Or the reader may wish to prove the existence of a fully labelled simplex on the interior by appealing directly to Sperner's lemma. The key idea is to *dualize* the simplex S to form a new simplex S^* . Loosely speaking, the dual of a simplex reverses the dimensions of k -dimensional and $(n - 1 - k)$ -dimensional faces. For instance, the corner vertices of S become the facets of S^* , and the facet barycenters of S become the vertices of S^* ; see Figure 10.




Figure 10. The dualization S^* of S . Vertices, barycenters, and one elementary simplex are marked to show how they are transformed.

A rigorous treatment of dualization can be found in Vick [15]. Note that S^* can be triangulated—in fact, using barycentric subdivision, the vertices and elementary simplices of S^* are in 1 - 1 correspondence with the vertices and elementary simplices of S . Let the triangulation of S^* inherit a labelling via this correspondence with S . One may now verify that the labelling of S^* is a Sperner labelling! Hence there exists a fully labelled elementary simplex of S^* , which corresponds to a fully labelled elementary simplex of S , as desired. This “dual” Sperner lemma is due to Scarf [16].

A constructive algorithm is obtained by following “trap-doors” in Sperner’s lemma. Choose an ϵ smaller than the rental difference for which housemates wouldn’t care (a penny?). Following trap-doors corresponds to suggesting pricing schemes and then asking various players, “Which piece would you choose if the rooms were priced like this?” Once a fully labelled elementary simplex is found, any point inside it corresponds to an ϵ -approximate rent-partitioning. We invite the reader to code a trap-door algorithm that could be implemented on a computer, one that would propose the necessary cut-sets and question the appropriate players at each step.

It is possible to obtain the Rental Harmony Theorem without any dualization argument and without condition (2) if one allows the possibility of negative rents. Specifically, let each person contribute a fixed amount K to a pool from which the

rent is paid. The leftover money is used to pay “rebates” associated with each room (which may be larger than K). This converts the problem into a fair division of goods (rebates), in which the space of rebates is a simplex that assumes a Sperner labelling if players demand a non-zero rebate. For large K this is quite reasonable. However, solutions may include situations in which a housemate is being *paid* by the others to live there. Thus allowing this possibility may not be realistic because in real life, paying housemates are more likely to ditch the subsidized housemate and use the extra room (and extra money) in other ways.

8. COMMENTS AND DISCUSSION. The Rental Harmony Theorem establishes the existence of envy-free chore division and a new ϵ -approximate algorithm, by simply thinking of the rent payments as chores and ignoring the rooms; divisibility of chores can be achieved by dividing the time spent on them. When reinterpreted, the three conditions from the Rental Harmony Theorem become: (1) all the chores must be assigned, (2) each person prefers no chores to some chores, and (3) preference sets are closed. These are pretty reasonable assumptions. The ϵ -approximate algorithm that arises from this does not involve a lot of cutting and reassembling, as do the exact methods proposed in [3] and [10].

For rent-partitioning, we point out that condition (2) may not always be a reasonable assumption. For instance, someone may be willing to pay a little bit of money for a room that is slightly larger than a free room. However, by inspecting the proof, one sees that the Rental Harmony Theorem still holds with a weakened version of condition (2):

Condition (2'). Each person never chooses the most expensive room if there is a free room available. This does not require the person to choose the free room.

In particular, this will hold if a person always prefers a free room to a room costing at least $1/(n-1)$ of the total rent. Hence condition (2') is a slightly weaker sufficient condition than that given by [1, pp. 1031–32]. To see why the Rental Harmony Theorem still holds, consider its proof and note that using this condition gives a more complicated labelling of S , but the corresponding labelling on S^* still remains Sperner.

What condition (2') does not address is a situation in which the total rent is so low, or some room so large, that one would be willing to pay for the most expensive room even when some other room is free. In practice, however, housemates do not usually choose a house with such lopsided arrangements. Even still, condition (2') can likely be weakened further, but the extent to which it can (and still maintain non-negative rents) is an open question.

Other triangulations may be used instead of barycentric subdivision. These have better convergence properties but are harder to describe; see [17] for a survey and applications to fixed point algorithms.

9. ANECDOTE AND ACKNOWLEDGMENTS. My first exposure to the Sperner argument for cake-cutting came via Michael Starbird, who attributed the method to a graduate student of his, Forest Simmons. Simmons had been presenting this cake-cutting scheme to math clubs and high school groups, but never formally submitted the idea for publication. His inspiration was the MONTHLY article by Stromquist [14], which made use of a theorem that can be proved by Sperner's lemma.

Many years later, when my friend Brad Mann told me about the rent-partitioning dilemma that he and his housemates were facing, I was reminded of these ideas and realized that Sperner's lemma could also be adapted to treat rent-partitioning, as well as chore division.

I am grateful to Arthur Benjamin, Steven Brams, Brad Mann, Forest Simmons, and Ravi Vakil for many helpful discussions, and I thank Michael Starbird for introducing me to Sperner's lemma.

REFERENCES

1. A. Alkan, G. Demange, and D. Gale, Fair allocation of indivisible goods and criteria of justice, *Econometrica* **59** (1991) 1023–1039.
2. S. J. Brams and A. D. Taylor, *An envy-free cake division algorithm*, Economic Research Reports, C. V. Starr Center for Applied Economics, New York University, 1992.
3. S. J. Brams and A. D. Taylor, *Fair Division: from Cake-Cutting to Dispute Resolution*, Cambridge University Press, Cambridge, 1996.
4. S. J. Brams and A. D. Taylor, An envy-free cake division protocol, *Amer. Math. Monthly* **102** (1995) 9–18.
5. D. I. A. Cohen, On the Sperner lemma, *J. Combin. Theory* **2** (1967) 585–587.
6. M. Gardner, *aha! Insight*. W. F. Freeman and Co., New York, 1978.
7. B. Knaster, C. Kuratowski, and S. Mazurkiewicz, Ein Beweis des Fixpunktsatzes für n -dimensionale Simplexe, *Fund. Math.* **14** (1929) 132–137.
8. H. W. Kuhn, Simplicial Approximation of Fixed Points, *Proc. Nat. Acad. Sci. U.S.A.* **61** (1968) 1238–1242.
9. R. Oskui, Dirty work problem, preprint.
10. E. Peterson and F. E. Su, Exact procedures for envy-free chore division, preprint.
11. J. M. Robertson and W. A. Webb, *Cake-Cutting Algorithms: Be Fair If You Can*, A K Peters Ltd., Natick, Mass., 1998.
12. F. W. Simmons, private communication to Michael Starbird, 1980.
13. E. Sperner, Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes, *Abh. Math. Sem. Univ. Hamburg* **6** (1928) 265–272.
14. W. Stromquist, How to cut a cake fairly, *Amer. Math. Monthly* **87** (1980) 640–644.
15. J. W. Vick, *Homology Theory*. Springer-Verlag, New York, 1994.
16. H. Scarf, *The Computation of Equilibrium Prices: An Exposition*, Cowles Foundation Discussion Paper No. 473, Cowles Foundation for Research in Economics at Yale University, November, 1977.
17. M. J. Todd, *The Computation of Fixed Points and Applications*, Lecture Notes in Economic and Mathematical Systems, New York, 1976.

FRANCIS SU earned his Ph.D. from Harvard in 1995 studying random walks, but his interest in fair division was a result of thesis procrastination. His research is now (fairly? randomly?) divided between studies of random walk convergence and fair division algorithms. He is an assistant professor at Harvey Mudd College, a “blue dot” MAA Project NExT Fellow, and an amateur songwriter in his spare time. Since the writing of this article he and HMC undergraduate Elisha Peterson wrote a web applet that uses an improved trapdoor algorithm; their *Fair Division Calculator* works for cakes, chores, or rent and they invite you to try it out at <http://www.math.hmc.edu/~su/fairdivision/>.
Harvey Mudd College, Claremont, CA 91711
su@math.hmc.edu