


Cutting Cakes Carefully

Theodore P. Hill and Kent E. Morrison

doi:10.4169/074683410X510272



Ted Hill is professor emeritus of mathematics at Georgia Tech, and has held visiting appointments in Costa Rica, Germany (Gauss Professor), Holland (NSF-NATO Fellow), Israel, Italy, and Mexico. He studied at West Point (B.S.), Stanford (M.S.), Göttingen (Fulbright Scholar), and Berkeley (M.A., Ph.D.), and his primary research interests are in mathematical probability, especially optimal-stopping theory, fair-division problems, and Benford's Law.



Kent E. Morrison is professor emeritus at California Polytechnic State University in San Luis Obispo, where he taught for thirty years. He has also taught at Utah State University, Haverford College, and the University of California at Santa Cruz, where he received his Ph.D. and B.A. degrees. Currently he is a visiting researcher at the American Institute of Mathematics in Palo Alto. He has a number of research interests in the areas of algebra, geometry, probability, and combinatorics.

People have argued about the distribution of food, and inheritances, and conquered territories since the dawn of history, but the first mathematical study of fair division seems to have been in the 1940s by renowned Polish mathematician Hugo Steinhaus. Since then, research in fair division has spread beyond mathematics to many fields including computer science, economics, law, and political science. Although the actual resource to be distributed may be telecommunication bandwidths, or land, or spoils of war, the underlying ideas in allocation procedures are most often described in the picturesque terms coined by Steinhaus, namely, the fair division of a cake.

In a typical model the “cake” must be divided among n people so that each person receives a *proportional share*, that is, a portion he values at least $1/n$ of his total value of the cake. Proportionality is a fundamental aspect of the mathematical theory of fairness.

Many essential and interesting issues immediately come to mind. Are there restrictions on how the cake can be cut? Do the players' values vary continuously? Are some fair divisions better than others? How does a player value the union of two different pieces? Does it pay to lie to other players? The main purpose of this article is to demonstrate, with a series of concrete examples, how seemingly minor changes in definitions and hypotheses can significantly alter the conclusions of fair-division theorems.

A good place to start is with one of the most elegant and practical of cake-cutting algorithms—Dubins and Spanier's classical moving-knife procedure [6]:

A knife is slowly moved at constant speed parallel to itself over the top of the cake. At each instant the knife is poised so that it could cut a unique slice of the cake. As time goes by the potential slice increases monotonically from nothing until it becomes the entire cake. The first person to indicate satisfaction with

the slice then determined by the position of the knife receives that slice and is eliminated from further distribution of the cake. (If two or more participants simultaneously indicate satisfaction with the slice, it is given to any of them.) The process is repeated with the other $n - 1$ participants and with what remains of the cake.

The only implicit assumptions are that values are nonnegative, additive, and, in the direction the knife is moving, continuous. The knife need not even be perfectly straight; it need not be moved perfectly parallel to itself; and the cake need not be simply connected (as traditional angel-food cakes, with a hole through the center, are not), nor even be connected (the cake could have been dropped on the table and landed in pieces). Clearly a player may take a chance for a larger share by not saying “stop” when he thinks the portion is fair, but only at the risk of receiving a sub-fair portion. Even though written informally, the ideas and conclusions in the moving-knife algorithm are clear and correct. However, as new fair-division procedures and criteria are sought, subtle changes to these classical hypotheses are sometimes crucial.

Cakes and values

The object to be divided will be called the *cake*, here denoted X . The *allowable partitions* of the cake consist of a non-empty collection of subsets of X , here denoted \mathcal{A} , which is assumed to include the empty portion \emptyset and whole cake X , and is assumed to be closed under complements and countable unions. In probability terms, \mathcal{A} is simply the set of “events.” In many situations such as the moving-knife procedure and most of the examples below \mathcal{A} is implicit from the context.

The value that a player gives to the various parts of the cake is represented by a measure on \mathcal{A} . Such a measure is a function $v : \mathcal{A} \rightarrow [0, 1]$ satisfying $v(\emptyset) = 0$ and $v(X) = 1$, so that every player regards the empty portion as having value zero and the whole cake as having value 1, or 100% of the total value. In addition, the value of the union of two disjoint pieces is the sum of the values of each. Thus, in probability terms again, v is simply a probability measure, but in cake-cutting, the number $v(A)$, rather than denoting the probability of the subset A , represents the relative value of the portion $A \in \mathcal{A}$ of the cake X to a player with measure v .

Does it matter how continuous the measures are?

Continuity of values may be interpreted in different ways. One interpretation is that the values are *atomless*, which means that no single point has positive value. A stronger interpretation is that each player’s value v is *absolutely continuous*, that is, given by the Riemann integral of a function f via the relationship $v(A) = \int_A f(x) dx$, where $f : X \rightarrow \mathbf{R}$, $f \geq 0$, and $\int_X f(x) dx = 1$. Here, we refer to this stronger interpretation simply as *continuity*.

The first example shows that if measures are atomless, but *not* continuous, then the moving knife procedure may not work.

Example 1. The cake X is the unit square and there are n players. The players’ measures are all concentrated on the top $T = \{(x, 1) : 0 \leq x \leq 1\}$, where the frosting is, and on T all the measures are uniform, that is, $v(\{(x, 1) : 0 \leq x \leq a\}) = a$ for all $0 \leq a \leq 1$, and $v(X \setminus T) = 0$. Those measures are atomless, but do not have p.d.f.’s on X , since $v(T) = 1$, but for any p.d.f. f , $\int_T f(x, y) dx dy = 0$. The moving-knife

procedure, if applied with the knife horizontal and moving from top to bottom, does not provide a fair division since every player says “stop” immediately. Moving the knife at any other angle, however, does guarantee an allocation where every player gets a portion valued at exactly $1/n$.

If a player’s measure is continuous, on the other hand, then every portion of an n -dimensional cake that has no n -dimensional volume, such as the top T of the cake in Example 1, must have measure zero. In particular, continuous measures are always atomless, but, as this example shows, not all atomless measures are continuous. If the measures are not only atomless, but are also continuous, then stronger conclusions often follow. For example, if the measures are continuous, then it is easy to see that the moving-knife method provides a fair division at *every* angle of the knife, in contrast to the situation in Example 1.

Does it matter whether measures are positive everywhere?

A continuous measure has a p.d.f. that by definition is non-negative, but some cake-cutting theorems require more. Some procedures require that the p.d.f.’s are positive everywhere, which implies that every player must like, to some extent, every part of the cake. On the surface this appears almost the same as simple non-negativity, but in this section we will see the important distinction between these two assumptions.

As a first example consider two procedures, both called *Surplus Procedure*, for dividing the unit interval between two players: SP [1], [2] and SP’ [4]. Both procedures require the two players to report their measures to a referee, who then determines the 50–50 points (medians). Both are supposed to *guarantee* that if the players’ 50–50 points are unequal, then each player will receive a portion he values strictly more than $\frac{1}{2}$. A third procedure, the *Divide-and-Conquer* (DC) algorithm for n players [3], is designed to give every player *strictly* more than a fair share, if their referee-computed medians are not equal.

However, without some additional assumption, such as that measures are positive everywhere, medians may not be unique. In that case, as the next example shows, all three procedures, SP, SP’, and DC, may fail to give every player a portion he values strictly more than $1/n$.

Example 2. The cake is the unit interval $[0, 1]$ and there are two players. Both players give zero measure to the middle interval $(\frac{1}{3}, \frac{2}{3})$ and uniform measure to the rest. That is, the p.d.f. of both players is $f(x) = \frac{3}{2}$ for $x \in [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$ and zero otherwise. Then Player 1’s 50–50 point could be marked at $\frac{1}{3}$ and Player 2’s at $\frac{2}{3}$, but since the measures are identical, neither can receive more than half without the other receiving less.

Efficiency and positivity Often more than one fair allocation of allowable portions to players is possible, and in that case it is natural to compare one given fair allocation with the other possible fair allocations. If no better allocations exist, the given allocation is called *Pareto optimal* or *efficient*. There are two interpretations of “better.” One is *weak Pareto optimality*, for which there is no other allocation that gives all players more of the cake. The second is *strong Pareto optimality*, for which there is no other allocation that is better for one person and at least as good for the others.

In [1], [2], and [4], the classical, two-player fair division procedure *cut-and-choose*, in which one player cuts and the other chooses, is claimed to be strong Pareto optimal.

However, the proof requires an additional assumption that the measures are positive everywhere. As the next example shows, without this assumption, cut-and-choose is not Pareto optimal.

Example 3. The cake X is the unit square $\{(x, y) : 0 \leq x \leq 1, 0 \leq y \leq 1\}$; there are two players; and \mathcal{A} is the collection of all pieces cut by a single straight line, i.e., the intersections of X with half-planes. Player 1 values only the top half of the cake and Player 2 only the bottom half; on those portions their measures are uniformly distributed. That is, the p.d.f. of Player 1 is $f_1(x, y) = 2$ if $\frac{1}{2} \leq y \leq 1$, and zero otherwise; and that of Player 2 is $f_2(x, y) = 2$ if $0 \leq y \leq \frac{1}{2}$, and zero otherwise. Clearly both measures are continuous, and the cutter, not knowing the other player's values, will always divide the cake into two parts he values equally. Thus the cutter, no matter how they slice will always wind up with exactly 50% of the cake. Observe, though, that the cut at $y = \frac{1}{2}$ will give each player a piece he or she values at 100% (see Figure 1).

For example, if Player 1 is the cutter and cuts vertically, then his unique optimal cut-and-choose solution is to bisect the cake along the line $x = \frac{1}{2}$, in which case each player receives a portion he values exactly at 50%. In this case, cut-and-choose is not even weak Pareto optimal. On the other hand, if Player 1 cuts horizontally, his unique optimal cut-and-choose point is the line $y = \frac{3}{4}$. He will receive a portion he values at exactly 50%, but then Player 2 chooses the bottom portion which he values at 100% of the cake. As before, an allocation of the top half of the cake to Player 1, and the bottom half to Player 2 is at least as good for Player 2, and is strictly better for Player 1. Thus if Player 1 cuts horizontally, cut-and-choose is weak Pareto optimal, but not strong Pareto optimal.




Figure 1. Cut 3 is better than Cut 1 for both players; Cut 3 is better than Cut 2 for Player 1, but not for Player 2. Cut 3 is a strong Pareto optimum; Cut 2 is a weak optimum, but not strong; Cut 1 is envy-free, but not Pareto optimal in either sense.

Envy-free allocation and positivity In [1], [2], and [5], it is claimed that, using $n - 1$ cuts, an envy-free allocation (an allocation in which no player values any other player's portion more than he values his own) is always strong Pareto optimal. The next example shows that the proof of this statement also requires adding an additional hypothesis, such as that measures are positive everywhere.

Example 4. The cake X is the unit interval $[0, 1]$ and there are two players. Player 1 values the cake uniformly and Player 2 values $[0, \frac{1}{4}]$ and $[\frac{3}{4}, 1]$ twice as much as




Figure 2. Left: Player 2 cuts as indicated, and player 1 chooses either half. Right: Player 2 cuts as indicated and player 1 chooses $[1/4, 1]$.

Player 1 and puts no value on $[\frac{1}{4}, \frac{3}{4}]$ (see Figure 2). If Player 1 is the cutter, his unique cut-point is at $x = \frac{1}{2}$, and each player will receive a portion he values at exactly $\frac{1}{2}$. The allocation of the interval $[0, \frac{1}{4}]$ to Player 2 and the rest to Player 1, however, gives Player 1 a portion he values $\frac{3}{4}$, and Player 2 a portion he values $\frac{1}{2}$ again, and so cut-and-choose is not weak Pareto optimal, and hence not strong Pareto optimal, either.

Similarly, to conclude as in [5] that that Stromquist’s well known envy-free moving-knife procedure for three players [12] is strong Pareto optimal among allocations using $n - 1$ cuts, also requires additional assumptions such as measures that are positive everywhere. This is shown in the next example.

Example 5. X is the unit interval $[0, 1]$, and there are three players. Player 1 values X uniformly, and Players 2 and 3 both place uniform value on $(0.4, 0.6)$ and no value on the rest. The Stromquist 4-knife, envy-free procedure, moving from left to right, gives portion $(0, \frac{1}{3})$ to Player 1 and portions $(\frac{1}{3}, \frac{1}{2})$ and $(\frac{1}{2}, 1)$ to Players 2 and 3. (It does not matter who gets which.) Thus, the Stromquist procedure allocates a portion to Player 1 that he values $\frac{1}{3}$, and Player 2 and Player 3 each get a portion that each values $\frac{1}{2}$. But the allocation $(0, 0.4)$ to Player 1, $(0.4, 0.5)$ to Player 2 and the rest to Player 3 is strictly better for Player 3 and equally good for the other two, and so the $n - 1$ cut Stromquist procedure is not strong Pareto optimal.

Does it matter whether the pieces are connected?

Although many division procedures require players to receive contiguous pieces of the cake, it may turn out, if non-contiguous pieces are allowed, that allocations using only contiguous pieces will fail to be Pareto optimal. Or, as stated in [5], “satisfying contiguity may be inconsistent with satisfying efficiency.”

In SP for example, the two players report their measures v_1 and v_2 to a referee, who determines the medians, a for Player 1 and b for Player 2. Without loss of generality, assume $a \leq b$. If $a = b$, the cake is cut at that point, and Player 1 receives the portion $[0, a]$ and Player 2 the rest. If $a < b$ the referee finds $c \in (a, b)$ so that

$$\frac{v_1(a, c)}{v_1(a, b)} = \frac{v_2(c, b)}{v_2(a, b)}.$$

The cake is cut at c and Player 1 receives $[0, c]$ and Player 2 the rest. If only allocations consisting of contiguous pieces are allowed and the values are positive everywhere it can be shown that SP is strong Pareto optimal. However, the next example shows that if non-contiguous pieces are allowed, then SP may not be Pareto optimal.

Example 6. The cake X is the unit interval $[0, 1]$. The set \mathcal{A} of allowable partitions is the collection of finite unions of subintervals of X , and there are two players with continuous values. Player 1's p.d.f. is 1.6 on $(0, \frac{1}{4})$ and on $(\frac{1}{2}, \frac{3}{4})$ and is 0.4 elsewhere; and Player 2's is 1.6 on $(\frac{1}{4}, \frac{1}{2})$ and $(\frac{3}{4}, 1)$ and 0.4 elsewhere. Then the procedure SP cuts the cake at $\frac{1}{2}$, and each player receives a portion worth exactly 0.5. But allocating $(0, \frac{1}{4})$ and $(\frac{1}{2}, \frac{3}{4})$ to Player 1, and the rest to Player 2, gives each player a portion he values 0.8, which is strictly better for each player. Thus SP is not even weak Pareto optimal.

The *Equitability Procedure* (EP) [1], [2], [4] for dividing the unit interval among $n \geq 2$ players has a similar dependency on the assumption of contiguous pieces. In EP, the players again report their values to a referee, who then determines cut-points $0 \leq c_1 \leq c_2 \leq \dots \leq c_{n-1} \leq 1$, so that $v_1(0, c_1) = v_2(c_1, c_2) = \dots = v_n(c_{n-1}, 1)$. The referee also determines analogous cut-points for the rest of the $n!$ orderings of the players, and then selects the ordering and cut-points for which this common value is greatest. Then $n - 1$ cuts are made at those points, and the intervals are allocated in that order. The claims in [1], [2], [4] that EP is strong Pareto optimal, however, may also fail without additional assumptions that the measures are positive everywhere and that contiguous portions are required. Without such assumptions, as the next example shows, EP is not in general even weak Pareto optimal.

Example 7. The cake X is the unit interval $[0, 1]$, and the set \mathcal{A} of allowable partitions is the collection of finite unions of subintervals of X . There are three players with measures that are positive everywhere. Player 1's p.d.f. is 2.4 on $(0, \frac{1}{6})$ and on $(\frac{1}{2}, \frac{2}{3})$ and is 0.3 elsewhere; Player 2's p.d.f. is 2.4 on $(\frac{1}{6}, \frac{1}{3})$ and $(\frac{2}{3}, \frac{5}{6})$ and 0.3 elsewhere; and Player 3's is 2.4 on $(\frac{1}{3}, \frac{1}{2})$ and $(\frac{5}{6}, 1)$ and 0.3 elsewhere. Then one EP solution allocates the portion $(0, \frac{1}{3})$ to Player 1, $(\frac{1}{3}, \frac{2}{3})$ to Player 3, and the rest to Player 2, in which case each player receives a portion worth exactly 0.45. But allocating $(0, \frac{1}{6})$ and $(\frac{1}{2}, \frac{2}{3})$ to Player 1, $(\frac{1}{6}, \frac{1}{3})$ and $(\frac{2}{3}, \frac{5}{6})$ to Player 2, and the rest to Player 3 gives each player a portion he values 0.8, which is strictly better for each player than the EP allocation, so EP is not weak Pareto optimal.

Does it pay to lie?

In fair-division theory, there are also complex and somewhat game-theoretic issues regarding what information is available to the players, whether the referee (if one is required) must act in a prescribed manner, and whether players must act truthfully. These, too, have been a source of confusion. For instance, there are several different notions of what it means for a procedure to provide an incentive for players to act truthfully.

In [1], [2], and [4] for example, an allocation procedure is defined to be *strategy-vulnerable* if a "player can, by misrepresenting its value function, *assuredly* do better, *whatever* the value function of the other player" (emphasis added). Otherwise, the procedure is called *strategy-proof*. However, it is easy to see that under this definition, every allocation procedure is strategy-proof, since if the players' measures all happen to be identical, assuredly they cannot *all* do better than a fair share, whether they act truthfully or not. In [9] strategy-vulnerable is re-defined to mean that a player can do at least as well and *sometimes* better by lying about his value measure. Under this new definition, it is claimed that SP is strategy-proof. However, the next example, adapted

from an example of K. Teh, shows that using SP, a player may sometimes gain a *risk-free* advantage by lying, so SP is not strategy-proof in this new sense.

Example 8. The cake X is the unit interval $[0, 1]$, and there are two players whose values are positive everywhere. Let Player 1's true measure have cumulative distribution function $F(x) = v_1([0, x])$. Since F is strictly increasing on $[0, 1]$, F has a unique median, say at $a \in (0, 1)$. Suppose that, rather than reporting his true values, Player 1 reports the c.d.f. F^* given by

$$F^*(x) = \begin{cases} \frac{1}{2} - 2\left(\frac{1}{2} - F(x)\right)^2 & \text{if } x \leq a \\ \frac{1}{2} + 2\left(\frac{1}{2} - F(x)\right)^2 & \text{if } x \geq a \end{cases}.$$

Now if Player 2 reports the same median a as Player 1, the SP cut-point is at a , and both players receive exactly half the cake. However, if Player 2 reports a median $b > a$, then Player 1 will receive a strictly larger portion by reporting the false F^* , instead of his true F , if and only if the cut-point c in (a, b) satisfies

$$\frac{F^*(c) - 1/2}{F^*(b) - 1/2} < \frac{F(c) - 1/2}{F(b) - 1/2}.$$

But this follows by definition of F^* since

$$\frac{F^*(c) - 1/2}{F^*(b) - 1/2} = \frac{2(F(c) - 1/2)^2}{2(F(b) - 1/2)^2} < \frac{F(c) - 1/2}{F(b) - 1/2},$$

since $c < b$ and F is strictly increasing. Figure 3 illustrates this for F uniform.




Figure 3. Left: The truthful p.d.f. f_1 of player 1 is uniform; the untruthful p.d.f. f_1^* devalues the interval around his median at $1/2$. Right: By lying player 1 gets more, since the cutpoint moves from c to c^* .

The argument for $b < a$ is similar, but in this case Player 1 receives the rightmost interval and the cut point c moves to the left giving Player 1 an even greater portion. Thus no matter what measure Player 2 reports, if Player 1 falsely reports F^* , then Player 1 will do at least as well and sometimes better than he would have done had he reported his true value. Therefore, SP is not strategy-proof in the sense of [9].

Of course, if Player 2 lies even more than Player 1 does, then Player 1 may receive less than he would receive if both players report their true measures. Unlike the

problems with Pareto optimality described above that can often be repaired by making additional assumptions such as positive values everywhere, there appears to be no quick fix to this problem with strategy-vulnerability.

Final remarks This article is intended to motivate others to discover the fascinating world of fair division. As these examples show, it is often possible to examine and test new and sometimes complex mathematical theories and hypotheses using basic ideas from algebra and calculus. Conversely, the classical cut-and-choose and moving-knife algorithms show it is possible to express practical, yet clean, clear, beautiful, and logical conclusions without highly technical language.

Acknowledgment. The authors are grateful to Walter Stromquist, two colleagues who wish to remain anonymous, four reviewers and the Editor for several corrections and many helpful suggestions.

Summary. This paper surveys the fascinating mathematics of fair division, and provides a suite of examples using basic ideas from algebra, calculus, and probability which can be used to examine and test new and sometimes complex mathematical theories and claims involving fair division. Conversely, the classical cut-and-choose and moving-knife algorithms show it is often possible to express practical, yet clean, clear, and beautiful logical conclusions without using highly technical language.

References

1. S. Brams, *Mathematics and Democracy*, Princeton University Press, Princeton NJ, 2008.
2. S. Brams, M. A. Jones, and C. Klamler, Better ways to cut a cake, *Notices Amer. Math. Soc.* **53** (2006) 1314–1321.
3. ———, Divide-and-conquer: a proportional, minimal-envy cake-cutting procedure, *Dagstuhl Seminar Proceedings* **07261** (2007); available at <http://drops.dagstuhl.de/opus/volltexte/2007/1221>, accessed 5/17/2010.
4. ———, Better ways to cut a cake—revisited, *Dagstuhl Seminar Proceedings* **07261** (2007); available at <http://drops.dagstuhl.de/opus/volltexte/2007/12217>, accessed 5/17/2010.
5. S. Brams and A. Taylor, *Fair-Division: From Cake-Cutting to Dispute Resolution*, Cambridge University Press, New York, 1996.
6. L. Dubins and E. Spanier, How to cut a cake fairly, *Amer. Math. Monthly* **68** (1961) 1–17. doi:10.2307/2311357
7. T. Hill, Mathematical devices for getting a fair share, *American Scientist* **88** (2000) 325–331.
8. T. Hill and K. Morrison, Counterexamples in the theory of fair division; available at <http://arxiv.org/abs/0901.2360>.
9. A. Magid, Editor's note: More on cake cutting, *Notices Amer. Math. Soc.* **55** (2008) 1072.
10. J. Robertson and W. Webb, *Cake Cutting Algorithms: Be Fair If You Can*, A K Peters, Wellesley MA, 1998.
11. H. Steinhaus, *Mathematical Snapshots*, 3rd ed., Dover, New York, 1999.
12. W. Stromquist, How to cut a cake fairly, *Amer. Math. Monthly* **87** (1980) 640–644. doi:10.2307/2320951

Compromise is the art of dividing a cake in such a way that everyone believes he has the biggest piece.

—Ludwig Erhard (1897–1977; Chancellor of West Germany, 1963–1966)

suggested by Julius Barbanel