9-1-2017

Carcassonne in the Classroom

Mindy Capaldi
mindy.capaldi@valpo.edu

Tiffany Kolba
tiffany.kolba@valpo.edu

Recommended Citation

Mindy Capaldi, & Tiffany Kolba. (2017). Carcassonne in the Classroom. The College Mathematics Journal, 48(4), 265. https://doi.org/10.4169/college.math.j.48.4.265

Carcassonne in the Classroom

Mindy Capaldi and Tiffany Kolba

Mindy Capaldi (mindy.capaldi@valpo.edu) received her Ph.D. in mathematics from North Carolina State University and is now an associate professor of mathematics and statistics at Valparaiso University. Although her dissertation topic was topology, she now studies mathematics education and the scholarship of teaching and learning. Capaldi's nonacademic interests include gardening, knitting, birding, and reading fiction. She would like to thank her coauthor for introducing her to Carcassonne (as well as many other great board games).

Tiffany Kolba (tiffany.kolba@valpo.edu) received her Ph.D. in mathematics from Duke University and is now an assistant professor of mathematics and statistics at Valparaiso University. Her research interests are in probability theory, and she especially enjoys probability applications to problems involving board games and twin birth rates.

"What is the probability of choosing a green ball from an urn with three blue balls, five green balls, and seven yellow balls?" Many students not only struggle to engage with this sort of question but are left wondering why the world of mathematics is obsessed with balls and urns. The variety of approaches to choose from makes probability a difficult subject for many students, yet probability is an important part of quantitative literacy since it is prevalent in everyday life. Finding ways to clarify probability for undergraduates is key to a successful mathematics experience.

One strategy for increasing student engagement with probability concepts is to teach probability through its application to games. Many previous works have investigated the use of Markov chains to model board games, such as Chutes and Ladders [**2, 4, 5**], Monopoly [**1, 3**], and Risk [**6, 7, 8**]. While these works have primarily focused on understanding the various games for their own sake, in this article we focus on using the board game Carcassonne in the classroom as a path for students to learn about probability through a more interesting context. In particular, we give a sequence of increasingly difficult probability problems derived from Carcassonne that can be used in a wide range of undergraduate mathematics courses.

Introduction to Carcassonne

Carcassonne is a tile-based board game first released in Germany in 2000. The name refers to a French town famous for its medieval fortifications. A year after its release, Carcassonne received the German Game of the Year and German Game Prize awards.

http://dx.doi.org/10.4169/college.math.j.48.4.265
MSC: 97A20, 97K50

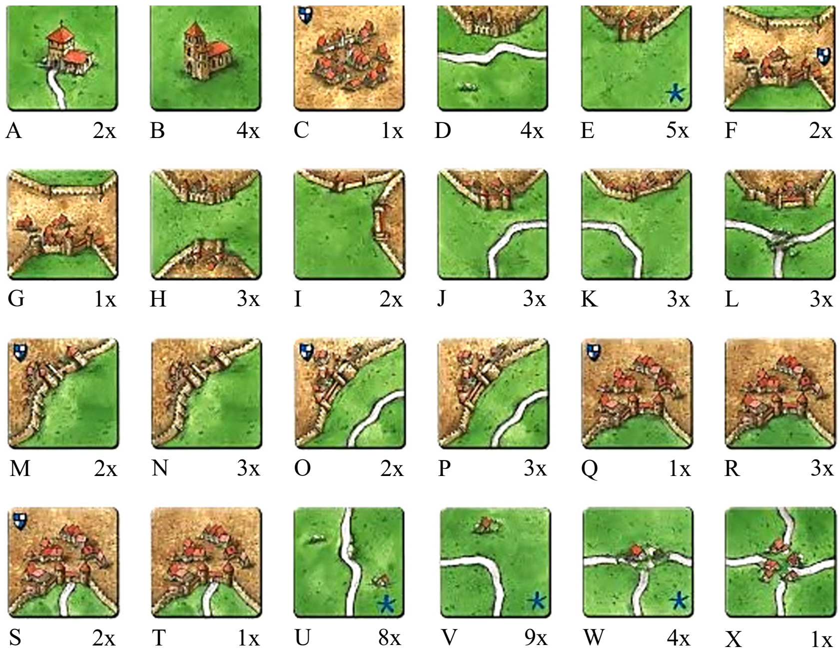

Figure 1. A list of the types and quantities of the tiles. The starting tile is of type D.

Its popularity spread to the United States along with many expansions of the game; there are now computer and mobile versions.

The object of Carcassonne is to build cities, roads, monasteries, and fields in order to gain points. Winning requires both luck and strategy. To begin the game, players randomly draw from 71 face down tiles and then must add a tile to the starting tile in a way that makes all features match; you cannot connect a road and a field, for instance. Each player has seven "meeple" pieces, which look like little people. Immediately after laying a tile, players choose whether or not to place their meeple on a feature (road, city, monastery, or field) of their tile. Some city tiles have a coat of arms, which double the tiles' points. Placing the meeple leads to points, depending on which feature is chosen. Figure 1 shows all game tiles along with their quantities and Table 1 gives point values. See `https://www.zmangames.com/en/products/carcassonne/` for a link to the Carcassonne rulebook.

While the game can seem complicated when first learning to play, students catch on fairly easily. (To save on time and difficulty, the field aspect of the game can be left out.) In a 50 minute class period, repeated in different semesters, students in a finite mathematics course were able to learn the rules and play the game at least once while also completing some probability problems.

Elementary problems

Carcassonne lends itself to basic probability problems that can be used to practice counting formulas or slightly more difficult topics like expected value. Courses such as finite mathematics, precalculus, or elementary statistics would be appropriate settings for the problems in this section.

One of the first probability formulas that students often encounter is $P(E) = |E|/|S|$ where E is an event in a sample space S that has equally likely outcomes and $|\cdot|$ denotes the size of a set. This is the essential formula for the following examples.

Problem 1. The monastery is a key type of tile because you can get up to nine points from this one tile. What is the probability of choosing a monastery if no tiles have been chosen yet?

Solution: Figure 1 lists the number of each tile type included, so we know that there are six monastery tiles out of 71 total tiles, not including the starting tile. Thus, the probability of choosing a monastery is $6/71$.

Here are several similar problems (given without solutions).

• What is the probability of choosing a monastery if 20 tiles have been picked and two of them were monasteries?

• At the point in the game described in the last question, what are the odds of choosing a tile that is not a monastery? (This emphasizes that odds are not a probability since the odds ratio is greater than one.)

• Suppose all of the coat of arms pieces have been picked as well as two tiles with just a straight road and four tiles with just a curvy road. What is the probability of choosing a tile that has a road or a monastery? What about a road and a monastery? (This emphasizes unions and intersections.)

Table 1. Point values for features in Carcassonne.

feature	points
completed city	2 per tile + 2 per coat of arms
unfinished city	1 per tile + 1 per coat of arms
road	1 per tile
monastery	1 + 1 for each surrounding tile
fields	3 for each completed city in the field

One of the slightly more difficult basic probability concepts for students to grasp is expected value. Since Carcassonne is a points-based game (as detailed in Table 1), questions about expected value are readily available.

Problem 2. At a given point in the game, assume that the following tiles are left:

- 3 monasteries (worth 5 points each because any of the three could be placed next to 4 other tiles)
- 21 tiles with a city but no coat of arms (worth 2 points each)
- 4 tiles with a city and a coat of arms (worth 4 points each)
- 6 tiles with only roads (worth 1 point each)

What is the expected value of the points you will earn on your next (randomly chosen) tile? Note: Only consider the points on the individual next tile, not points from features to which it may connect.

Solution: The size of the sample space is 34, the number of remaining tiles. Summing over the product of the payoff and probability of each outcome gives

$$5\left(\frac{3}{34}\right)+2\left(\frac{21}{34}\right)+4\left(\frac{4}{34}\right)+6\left(\frac{1}{34}\right)=\frac{79}{34}.$$

The final elementary probability problem uses the multiplication principle and distinguishes between counting and probability.

Problem 3. At the beginning of the game, how many ways are there for the first four tiles that are picked to be a monastery, then a city with no roads, then a city with a road, and finally a road with no city or monastery? What is the probability that this happens?

Solution: There are six monasteries, 23 cities with no roads, 20 cities with roads, and 22 roads with no cities or monasteries. The number of ways to choose the given four tiles is $|E| = 6 \cdot 23 \cdot 20 \cdot 22 = 60720$ and the probability of this event happening in the first four turns is

$$P(E) = \frac{|E|}{|S|} = \frac{60720}{71 \cdot 70 \cdot 69 \cdot 68} = \frac{60720}{23319240} \approx .003.$$

Intermediate problems

This middle level of problems incorporates the more involved topics of Bayes's theorem and Markov chains. The first two problems can be easily customized, with the basic idea and level of difficulty remaining the same. They can also be made more challenging by considering situations where tiles have already been played.

We use Bayes's theorem relating conditional probabilities to investigate the field component of the game. A field is any area of connected grass not separated by a road or city. Fields can be claimed by placing a meeple, in this case a farmer, on the grass feature of one of those tiles as it is played. Field points can be game changers, often determining the winner, but using a meeple on a field is risky because once it is placed you do not get that meeple back.

Problem 4. Consider the following five mutually exclusive types of tiles: monasteries (E_1), cities with a coat of arms (E_2), cities with a road but no coat of arms (E_3), cities with no road or coat of arms (E_4), and tiles with only roads (E_5).

Suppose that in previous games you have monitored your opponent, MacDonald, and determined the likelihood that he will choose to place a farmer on a field (call this event F) when each of these tiles is drawn. MacDonald never chooses a field over a monastery. Coat of arms pieces are worth higher points, and one of these pieces does not have a field, so he only chooses to place a farmer 3 times out of 10 for type E_2. Sixty percent of the time he places a farmer on type E_3 and half the time on E_4. Finally, three-fourths of the time he chooses fields over simple roads.

If you start a new game with player MacDonald where he gets to play a tile first, what is the probability that a road (E_5) was drawn if he placed a farmer on the tile?

Solution: We can summarize the probabilities given above as the following conditional probabilities:

$$P(F|E_1) = 0,\ P(F|E_2) = \frac{3}{10},\ P(F|E_3) = \frac{6}{10},\ P(F|E_4) = \frac{1}{2},\ P(F|E_5) = \frac{3}{4}.$$

Working from Figure 1, the counts for the events are

$$|E_1| = 6,\ |E_2| = 10,\ |E_3| = 16,\ |E_4| = 17,\ |E_5| = 22,$$

not including the starting tile in E_3. The ratios of these numbers to the 71 total number of tiles gives the probability of MacDonald randomly picking any of the five types.

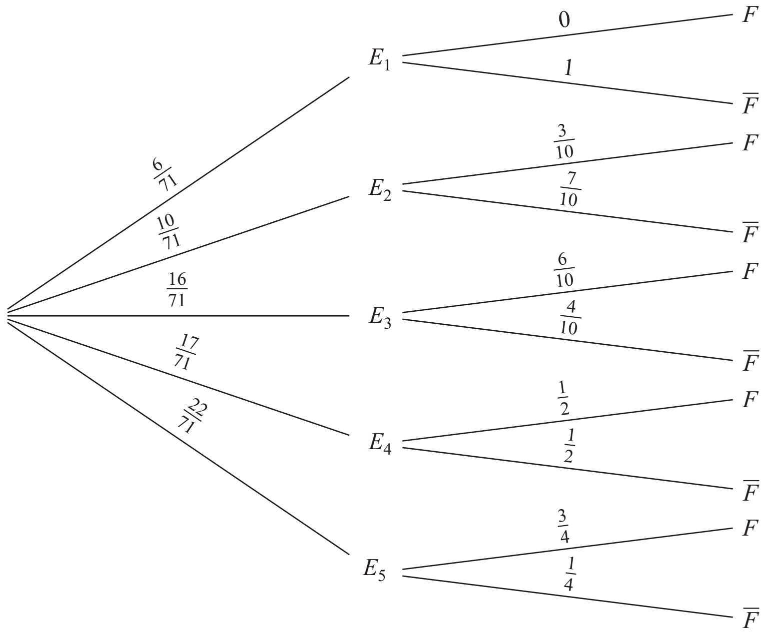

Figure 2. Tree diagram for Problem 4.

The information can be nicely organized as a tree diagram shown in Figure 4 where $\overline{F}$ denotes the complement of the event F, the event that he does not place a farmer on a field.

Recall that we want to determine $P(E_5|F)$. By Bayes's theorem,

$$P(E_5|F) = \frac{P(E_5)P(F|E_5)}{P(F)}.$$

The probabilities in the numerator can be found directly from the tree diagram: $P(E_5) = 22/71$ and $P(F|E_5) = 3/4$. To find $P(F)$, we use each probability that branches to F on the tree. That is,

$$\begin{aligned} P(F) &= P(E_1 \cap F) + P(E_2 \cap F) + P(E_3 \cap F) + P(E_4 \cap F) + P(E_5 \cap F) \\ &= P(E_1)P(F|E_1) + P(E_2)P(F|E_2) + P(E_3)P(F|E_3) + P(E_4)P(F|E_4) \\ &\qquad + P(E_5)P(F|E_5) \\ &= \frac{6}{71} \cdot 0 + \frac{10}{71} \cdot \frac{3}{10} + \frac{16}{71} \cdot \frac{6}{10} + \frac{17}{71} \cdot \frac{1}{2} + \frac{22}{71} \cdot \frac{3}{4} \approx .530. \end{aligned}$$

Finally, we use Bayes's theorem to find

$$P(E_5|F) = \frac{\frac{22}{71} \cdot \frac{3}{4}}{P(F)} \approx .439.$$

A follow-up question that builds off the work of the previous example involves the idea of independent events.

Problem 5. Are the events of MacDonald drawing a road tile (E_5) and playing a farmer (F) independent?

Solution: We already know that $P(E_5) = 22/71 \approx .310$ and $P(E_5|F) \approx .439$. Since these probabilities are not equal, the events are dependent.

Another intermediate problem concerns Markov chains, memoryless processes where the probability of transitioning to the next state depends only upon the current state and not the past. As mentioned earlier, Markov chains have been used to model many other board games, but Carcassonne does not easily fit a Markov chain model. When studying Markov chains, students are often given numerous examples that satisfy the definition, but are not usually presented with many situations which do not fit the criteria of a Markov chain. In other words, they are often not given many nonexamples, which can be as illuminating as examples.

Problem 6. There are 24 distinct types of tiles in Carcassonne, labeled A–X, as depicted in Figure 1. Let X_n denote the tile type of the nth randomly selected tile. (Note that $X_0 = \text{D}$.) Is X_n a Markov chain? Why or why not?

Solution: No, because the probability of the next tile type depends not just on the current tile type drawn but also on which tile types have been drawn in the past.

Here are two similar problems (given without solutions).

- Let X_n denote the number of unfinished cities after the nth randomly selected tile has been placed. (Notice that $X_0 = 1$ since the starting tile has a city.) Is X_n a Markov chain? Why or why not? (This could be modified to consider roads or monasteries instead.)
- Define a process related to Carcassonne that is not a Markov chain. Be sure to clearly define what X_n and n represent and articulate why it does not satisfy the definition of a Markov chain.

Challenging problems

One of the most interesting aspects of Carcassonne is the tile laying dynamics: Rather than having a fixed board as in many board games, the board in Carcassonne is dynamically built with a tile added each turn. When analyzing valid tile placements, students can explore the ideas of topological equivalence and equivalence classes by discovering that only the edges matter, i.e., whether each edge is a city, field, or road. What is printed on the interior of the tile (such as a monastery or coat of arms) is irrelevant. Students can also explore the effects of rotating a tile and discover that a tile with edges ordered [city, road, field, road] is equivalent to a tile with edges ordered [field, road, city, road] but not equivalent to a tile with edges ordered [city, road, road, field].

Problem 7. Figure 1 delineates 24 distinct tile types labeled A–X. When just considering tile laying dynamics, which tile types can be considered equivalent? How many different equivalence classes are there?

Solution: Equivalence classes are {F, G, H}, {I, M, N}, {O, P}, {Q, R}, {S, T}, and the 12 remaining types each as a singleton class, giving 17 total. (It is important for students to realize that J and K are not equivalent since the relative order of the edge types is different, even under rotation.)

An important aspect of designing a tile laying game is ensuring that there is a high probability that there is at least one valid location available to place each tile. The next example allows students to explore the likelihood that there are no available valid placements for a tile in Carcassonne.

Problem 8. Assuming that each tile is equally likely to be placed in any of the available valid locations and counting each distinct orientation as a distinct valid location, what is the probability that the first randomly selected tile cannot be validly placed? What about for the second randomly selected tile?

Solution: The first randomly selected tile has a 0% chance of no valid locations since the starting type D tile connects to any tile (i.e., it has city, field, and road edges). The analysis for the second tile is much more complex. We will see that the number of valid locations is zero only in two cases.

The first case is where the second tile selected has all city edges (type C) and the first tile selected is placed to complete the starting city with no remaining free city edges (so type D, E, J, K, or L). An example of this case is depicted in the picture on the left in Figure 3, where the first selected tile is type J. Tile types D, J, and K each have a 3/71 probability of being the first tile selected and each have six possible locations (counting different orientations in the same spot). Since these tiles have six possible locations, but only one location with a city edge, there is a 1/6 probability that the placement of one of these tiles would block the placement of a second tile of type C. Similarly, tile type E has a 5/71 probability and four possible locations and tile type L has a 3/71 probability with seven possible locations. Tile type C has a 1/70 probability of being the second tile selected. Therefore, the probability of the first case is

$$3 \cdot \frac{3}{71} \cdot \frac{1}{6} \cdot \frac{1}{70} + \frac{5}{71} \cdot \frac{1}{4} \cdot \frac{1}{70} + \frac{3}{71} \cdot \frac{1}{7} \cdot \frac{1}{70} \approx 0.00064.$$

Figure 3. Illustration of the two cases where there is no valid placement for the second tile selected.

The second case is when the second tile selected has all field edges (tile type B) and the first tile selected is placed on the field edge of the starting tile with no remaining field edges (so type D, J, K, Q, R, or W). An example of this case is depicted in the picture on the right in Figure 3, where the first tile selected is of type J. Tile types D, J, and K again each have a 3/71 probability of being the first tile selected and each have six possible locations (counting different orientations in the same spot). Tile types Q and R each have four possible locations and together have probability 4/71 of being the first tile selected. Tile type W has a 4/71 probability of being selected first with seven possible locations. Tile type B has a 4/70 probability of being the second tile selected. Therefore, the probability of the second case is

$$3 \cdot \frac{3}{71} \cdot \frac{1}{6} \cdot \frac{4}{70} + \frac{4}{71} \cdot \frac{1}{4} \cdot \frac{4}{70} + \frac{4}{71} \cdot \frac{1}{7} \cdot \frac{4}{70} \approx 0.00247.$$

Thus the total probability that the second tile selected has no valid possible locations is approximately $0.00064 + 0.00247 = 0.00311$.

Problem 8 can be modified to allow for a nonuniform distribution for the placement of the first tile among its possible locations. For example, in actual game play, if the first tile selected can complete the city on the starting tile, there is likely a high probability that the player would do so, which would then result in a higher probability that the second tile has no valid possible locations.

The analysis of the probability of no valid locations for the third randomly selected tile is even more complex than for the second tile. This increasing complexity can be used to motivate the need and usefulness of computers to perform simulations. In an upper-level probability course or computer science course where students are learning Monte Carlo methods, performing simulations related to aspects of the game of Carcassonne can help increase student interest in the topic.

Problem 9. Using Monte Carlo simulation, approximate the expected number of tiles that have to be discarded during a game of Carcassonne due to no valid placement locations. Assume that each tile is equally likely to be placed in any of the valid locations, counting each distinct orientation as a distinct location.

Solution: Monte Carlo simulation gives an expected number around 0.02 tiles. So on average, we would expect to have to discard one tile out of about every 50 games of Carcassonne due to no valid placement locations. (Note that this problem can be used to help students explore the subtleties of simulating rare events.)

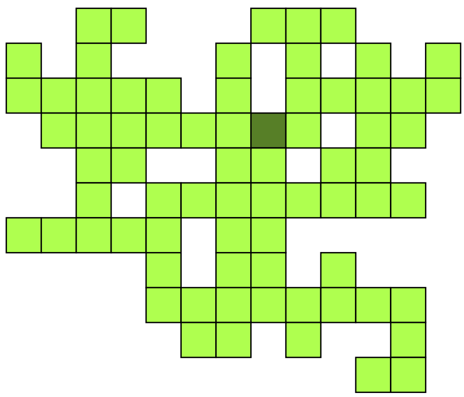

Figure 4. Final board layout from one simulation of the game of Carcassonne. The starting tile is indicated in dark green.

Another aspect of a tile-laying game such as Carcassonne that a game designer must consider is how the dimensions of the playing surface restrict where tiles can be placed. Figure 4 shows the layout of the final board configuration for one simulation of the game, assuming a uniform distribution for each tile among its possible locations, as well as no restrictions on the playing surface. The final dimensions are 13 tiles wide by 11 tiles tall, corresponding to a surface that needs to be at least 23-9/16” by 19-15/16” (each tile is 1-13/16” square). When actually playing Carcassonne, the table or desk size may prevent many tiles from being placed in otherwise valid locations. Students could repeat the simulation in Example 9 under different surface size restrictions and predict beforehand how the expected number of discarded tiles will increase as the dimensions of the playing surface decrease. The website `https://sites.google.com/a/valpo.edu/carcassonne/` contains annotated MATLAB code that simulates the tile laying dynamics. This code can be freely downloaded and used as a starting point for basic simulations that students can then build upon to simulate more detailed aspects of Carcassonne, such as where meeples are placed and tracking point values. The possibilities are virtually endless!

The authors hope that they have sparked interest in the use of Carcassonne for research and pedagogy, and welcome public use of the website https://sites.google.com/a/valpo.edu/carcassonne/ to aid others' investigations.

Summary. Learning about probability can be hard and frustrating for many students. However, learning about probability through examples with board games can make this task more interesting and fun. We present a sequence of increasingly difficult probability problems derived from the popular board game Carcassonne. Each question is appropriate either for a college classroom or for undergraduate research, with topics including basic counting problems, expected value, Bayes's theorem, Markov chains, and Monte Carlo simulation. Some problems have solutions, but other questions are left open for the reader to explore.

References

1. S. D. Abbott, M. Richey, Take a walk on the boardwalk, *College Math. J.* **28** (1997) 162–171, http://dx.doi.org/10.2307/2687519.
2. S. C. Althoen, L. King, K. Schilling, How long is a game of Snakes and Ladders?, *Math. Gaz.* **77** (1993) 71–76, http://dx.doi.org/10.2307/3619261.
3. R. B. Ash, R. L. Bishop, Monopoly as a Markov process, *Math. Mag.* **45** (1972) 26–29, http://dx.doi.org/10.2307/2688377.
4. L. A. Cheteyan, S. Hengeveld, M. A. Jones, Chutes and Ladders for the impatient, *College Math. J.* **42** (2011) 2–8, http://dx.doi.org/10.4169/college.math.j.42.1.002.
5. R. W. Johnson, Using games to teach Markov chains, *PRIMUS* **13** (2003) 337–348, http://dx.doi.org/10.1080/10511970308984067.
6. J. A. Osborne, Markov chains for the Risk board game revisited, *Math. Mag.* **76** (2003) 129–135, http://dx.doi.org/10.2307/3219306.
7. P. Pierce, R. Wooster, Conquer the world with Markov chains, *Math Horiz.* **22** no. 4 (2015) 18–21, http://dx.doi.org/10.4169/mathhorizons.22.4.18.
8. B. Tan, Markov chains and the Risk board game, *Math. Mag.* **70** (1997) 349–357, http://dx.doi.org/10.2307/2691171.