Beyond Representativeness: Productive Intuitions About Probability

Phil Vahey, Noel Enyedy and **Bernard Gifford**
Education in Mathematics, Science and Technology
4533 Tolman Hall
University of California, Berkeley
Berkeley, CA 94720-1670
pvahey@socrates.berkeley.edu, enyedy@socrates.berkeley.edu, and
bgifford@socrates.berkeley.edu

Abstract

Although research has found many flaws in people's probabilistic reasoning, we have found that middle-school students have many productive ideas about probability. This study examines the probabilistic reasoning used by middle-school students as they used a technology-mediated inquiry environment that was conceptualized and developed to engage students in the task of analyzing the fairness of games of chance. This research demonstrates that students employ productive probabilistic reasoning when participating in this task, and also demonstrates that commonly reported heuristics such as representativeness do not adequately describe student reasoning.

Prior Findings in Probabilistic Reasoning

There is a rich literature based on the many misconceptions people display when asked to reason probabilistically. By far the most influential work has been by Tversky & Kahneman (1982), who showed that much of people's probabilistic reasoning could be described by the heuristics of representativeness and availability. The *representativeness* heuristic is characterized by making judgments based on the degree to which *A* is representative of, or resembles, *B* (Tversky & Kahneman, 1982). This representativeness heuristic has been used to explain insensitivity to sample size, the gambler's fallacy, the base-rate fallacy, incorrect judgments about the output of random processes, and other non-normative judgments. The *availability* heuristic is characterized by making judgments based on the ease with which instances of a certain event can be brought to mind. This heuristic has been used to explain biases due to the retrievability of instances, biases due to the effectiveness of a search set, and biases of imaginability.

In addition, many other misconceptions have been identified. For example, people may believe that there is a lack of variability in the world, people have too much confidence in small samples, people do not see the importance of small differences in large samples, and people seem unaware of regression to the mean in their lives (for an overview of such misconceptions, see Shaughnessy, 1992).

Several attacks have recently been made on this literature. Roughly speaking, these criticisms come in two forms: (i) the so-called biases and misconceptions are not due to faulty probabilistic reasoning, but are due to situational factors from the experimental design, including the posing of intentionally misleading questions (Gigerenzer, 1996; Hilton, 1995; Konold, et al. 1993); and (ii) these heuristics do not explain the cognitive processes used by people when reasoning under uncertainty, and such heuristics were not derived from protocol analysis, but were inferred from questionnaire data (Gigerenzer, 1996; Konold, et al. 1993; Lajoie et. al, 1995). Our research situated seventh grade students in an environment where we hoped to minimize misleading situational factors, and we analyzed verbal data to allow us to investigate student reasoning about probability.

The Study

In this study pairs of seventh grade students collaborated in the Probability Inquiry Environment (PIE). PIE was created as a collaborative guided-inquiry environment (cf. White, 1993) in which students are asked to evaluate the fairness of games of chance. In PIE students use representations and tools such as event trees, simulations, and real-time graphs and histograms. PIE's inquiry cycle consists of being introduced to the games, making predictions about the games, running simulations of the game, and then drawing conclusions based on the simulation.

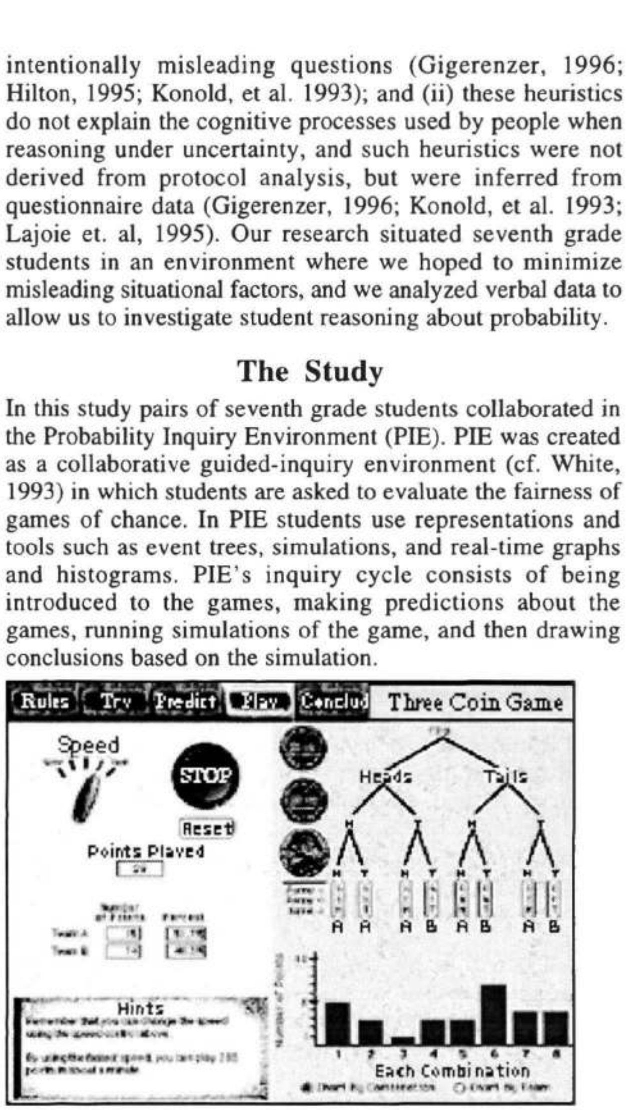

Figure 1: The PIE interface.

In this study students were asked to analyze two games of chance to determine if they were fair. The first game is called the Two-Penny game, where Team A scores a point whenever both coins come up the same (heads-heads or tails-tails), and Team B scores a point whenever both coins come up differently (heads-tails or tails-heads). This game is fair, as all outcomes are equally likely, and each team scores on two out of the four possible outcomes. The second game is called the Three-Coin game, where Team A scores a point on five of the eight possible outcomes, and Team B scores a point for three of the possible outcomes. Because each outcome is equally likely and Team A scores on more outcomes than Team B, this game is unfair in favor of Team A. An event tree that enumerated all the possible outcomes and that visually presented the scoring combinations for each team was on the screen at all times, as was a dynamic histogram that showed scoring either by each combination of coins, or by each team (see Figure 1).

Methodology

The research team recruited eight single-sex pairs of students from a local urban middle school. All of the sessions took place in the summer, between the students' sixth grade and seventh grade school years. Four of the pairs were boys, and four of the pairs were girls. The students represented a wide range of ethnicities. The students were paid $5 an hour, and spent about an hour and a half using PIE during one two hour session. During this session they were videotaped, and PIE recorded their actions on the computer. A researcher was always in the room with the students, and would occasionally interject to help clear up any confusion arising from the PIE interface.

The data used in this study consists of student discussions as they participated in the activities, and their responses to the on-line prediction and conclusion questions. The videotape data of the students using PIE were transcribed, and these transcripts were combined with the data recorded by PIE to create a record of all student discussions and student interactions with PIE during the session. These transcripts were then analyzed, and all instances of students' reasoning about the games was found. Although, by the end of the study, seven of the eight pairs of students were able to reason normatively about the games, this paper will not concentrate on the events that led to this normative reasoning. Instead, the students' reasoning throughout the entire session will be compared to the way in which normative probabilistic reasoning could have been employed in those situations.

First, to acquaint ourselves with an example of student reasoning in PIE, a case study of student reasoning during the session will be presented. This will be followed by a comparison between normative reasoning and student reasoning, in an attempt to better understand students' probabilistic reasoning. This new understanding of students' probabilistic reasoning will then be compared to the existing misconceptions literature to see if this understanding can account for the findings of others.

A Case Study of Q and T Using PIE

When Q and T felt that they understood the two-penny game (Figure 2), they made their predictions. When predicting that the Two-Penny game is fair, they explicitly assigned a 50% chance to the combinations of coins that score a point for each team, resulting in a final answer that is perfectly aligned with normative reasoning (Figure 3). Note however, that the students never explicitly justified this 50% chance, and we will not attempt to make claims about the students' probabilistic reasoning in this instance.

T: I get it, if they're both heads team A gets a point, if they're both tails team A gets a point, and if they're one heads and one tails Team B gets a point.

Figure 2: Understanding the two-penny game.

Q: We think the game is fair because you have a 50% chance of getting both heads and both tails.
Typed: We think that the game is fair because you have a 50% chance of getting both heads and both tails. You also have a 50% chance of getting one tails and one heads.

Figure 3: Fair, based on a 50% chance for each team.

Q and T next answered what they meant by fair. Although student ideas of fairness were interesting, most students, including Q and T, decided that a game is fair if all teams have an equal chance of winning, a normative view of fairness. Student ideas of fairness will not be further discussed in this paper.

In the next prediction question, Q and T were asked to manipulate histograms to make predictions about what would happen after 10 points, and after 200 points. T stated that heads would occur more, so she would expect Team A to win more. Q countered this by saying that coins usually come up differently (this is consistent with representativeness), so she would expect Team B to win more. They then decided that the game will most probably be tied (Figure 4). After further discussion, the students decided that luck would be an important factor in the game (Figure 5), and this meant that Team A might win sometimes, and Team B might win other times. Again the students provided final answers that were close to normative, but an analysis of the verbal data shows that their reasoning processes were quite rich and invoked more intuitions than are found in just their final answer.

T: OK, what about this, I think Team A will win. Because imagine all of the heads we are gonna get
Q: I think that Team B would win it, because when you throw it it's like real luck when you get both of them the same. When you throw it, most of the time they land differently. [pause] I think they'd probably be tied, but if I had to choose one, see, this looks right to me [pointing to even histogram bars]
T: That's what I think too

Figure 4: Who will win more?

T: I think that this is just a game of luck...this is like a game of guessing
Q: It's really like someone has to win, because it's like you win some you lose some, it's not like a permanent game.
[T moves histograms so A is winning after 10, and B is winning after 200]
Q: Now how can purple [Team A] be winning on this one and green [Team B] be winning on that?
T: Well, once you win you don't always win

Figure 5: A game of luck

For the final prediction question in the two-penny game, the students were asked if any of the combinations (outcomes) would happen more than any other. The students were again asked to manipulate a series of histograms, and type in a justification. Although Q and T started to reason about the different combinations, they then switched to talking about the probability of heads or tails. They decided that because this is not a game of skill, but a game of luck, they would expect heads and tails to come up the same amount (Figure 6).

Q: We don't think that any combination should happen more than any other, because it's luck
T: It's how you throw it
Q: Plus, it's a game of luck
T: I think this coin game, it's not a game of skill.
Q: No, no no. Don't, what are we writing?
T: We don't think that, it's not a game of skill
Q: That one penny will come up more than the other
T: It's all equal...wait a minute, we don't think that heads or tails is more likely to come up than the other. We don't think that heads or tails will come up more
Typed: We don't think that heads would come up more than tails. We think that this game is a game of luck.

Figure 6: Will any combination happen more?

Q and T then started the game simulation. After only four points they noted that tail-head was happening more than the other combinations. The game then stopped after ten points, and told the students that they could either look at the results or continue playing. At this point Q wanted to go back into predictions to see if the predictions agreed with the results (note that this behavior is consistent with the well-documented law of small numbers). The researcher asked them to continue playing, telling them that they would be able to modify their predictions later, so Q and T continued in Play.

After playing for several more points, Q and T then decided to run the simulation at the fastest speed. At this speed the game runs ten points at a time, and individual coin flips can not be perceived. After about 20 seconds the game reached 200 points, and gave the students the option to continue playing or stop and analyze the results. Q and T chose to stop playing and go immediately into Conclude.

The first three conclusion screens asked the students to evaluate their predictions. For each of these conclusions, the students stated that, although their predictions did not exactly match the actual data, the results were close enough for them to still agree with their predictions (Figure 7). Additionally, in Figure 8, Q came back to their earlier statement that being a game of luck is an important aspect of the game. This became the single most important factor for these students for the remainder of the session.

T: Even, look, almost even
Q: Yeah, so it's pretty fair because no one is like way, way more than the other

Figure 7: Conclusion—is the game fair?

Q: No, OK, why? We keep saying the same thing over. Like this evidence-no because this is just a game of luck. And they're all equal anyway.

Figure 8: Conclusion—are any combinations more likely?

For the last set of conclusions the students were asked if the number of combinations that score a point for each team is an important factor in determining fairness. However, Q and T understood this question to be asking about the data already collected. Although they first stated that this data is important, they then decided that luck is more important, and one doesn't need to know anything about the number of outcomes (Figure 9). Then, when asked to state the most important thing in determining if the game is fair, they again stated the importance of luck (Figure 10).

Q: Very important, don't you think...why is it very important?
T: It's data
Q: It's important data, and umm
T: You need the data to play the game
Q: You need to know the number of combinations...But it's not that important though, as a matter of fact, it's not important at all, cause it's a game of luck. Yeah, it's not important
T: It's a little important...
Q: ...but the game is just a game of luck anyway. So if you didn't have the data, it wouldn't matter anyway.
Typed: It is important because it's data, but on the other hand it is not that important because it's just a game of luck.

Figure 9: It's just a game of luck

Q: The most important thing is that you understand that the game is just luck

Figure 10: The most important thing in determining fairness

Q and T were then introduced to the Three-Coin game, and Q once more decided that the game was fair because it was a game of luck, although T was hesitant to agree. However, T could not state why she thought the game was unfair, and finally determined that, since all the outcomes were possible, the game must be fair. Note that, although the partitioning of the outcomes into points for A and points for B were on the screen at all times, the students never considered counting the outcomes to determine if each team had an equal number of outcomes. Although it is dangerous to make inferences based on the *absence* of an action, the fact that it never occurred to these (or most) students to simply count up the number of outcomes, especially when something about the game seemed troubling, may point to a lack of an understanding of the importance of the outcome space in determining probabilities.

Then, consistent with their predictions for the Two-Penny game, and consistent with their idea that luck means that a game is fair, for the remainder of the predictions Q and T stated that the teams will score an approximately equal number of points, and each of the combinations should occur equally. They then put the game on the fastest speed, and quickly played up to two hundred points. When the game reached two hundred points, they went into conclude and simply agreed with all their predictions, without comparing their predictions to the actual results, even stating that the data is not important (Figure 11). Such reasoning is based, presumably, on the statements made at the end of the Two-Penny game, that a game of luck must be fair, and data is not an important factor.

Typed: THE COINS IS JUST DATA AND THE DATA IN THIS GAME IS NOT THAT IMPORTANT.

Figure 11: Conclusion—are the combinations in the Three-Coin game important?

At this point the researcher asked them to play some more, reminding them of the Reset button that sets the points back to zero. After playing several more rounds up to two hundred points, T decided that, because Team B kept losing, the game must be unfair. Q, however, kept stating that the game is just luck, and so must be fair (Figure 12). T did not accept this answer, and continued looking for an explanation. She finally noticed the difference in the number of outcomes that scored a point for each team, and then resorted to a strategy of making the game fair, contrasting a partitioning of points that would make the game fair with the actual partitioning (Figure 13). Q then understood how this partitioning was relevant and agreed that the game was unfair.

T: See, Team B is losing by a lot. Told you it was unfair
...
Q: This game is just luck, it's just a penny game, it's just luck

Figure 12: Unfair versus luck

T: Why is it fair, Q?
Q: Because it's a game of luck, it's jut throwing pennies
T: I'm not talking about whose got the penny, I'm talking about right here [pointing to the tree]. They keep losing, and I'm trying to figure out why...wait a minute! See how this is AA right here? and this is AB AB AB
R: Mmm-hmmm
T: But shouldn't it be BB?
R: What do you mean, shouldn't it be?
T: Right here it says AA AB AB AB

Figure 13: Making the game fair

Summary: Q and T's analysis of the Two-Penny game began with them stating several different, often competing or conflicting, intuitions about probability, few of which seemed to carry any deep commitment. And, although representativeness could be used to describe some of their reasoning, it is a far from adequate account, as much of their reasoning is inconsistent with representativeness. Q and T then began to consider luck the single most important aspect of the game, even stating that they did not need data to determine the fairness of the games. Their commitment to this position was shown in the Three-Coin game, when Q explicitly denied the importance of data that showed that this game was unfair. Note that she did not fall prey to the law of small numbers, nor did she dismiss the game as "cheating", nor did she suffer from confirmation bias, misinterpreting the data as showing that the games were tied. Instead, she acknowledged the results of the simulation and simply said that these results were not relevant. It was not until T was able to determine that the partitioning of the outcome space was unequal that they were able to confirm that the game was unfair, and it took a notably long time until this counting strategy was employed by the students, lending credence to the supposition that the outcome space was not a salient feature of this situation. Note that most of this behavior is not consistent with the heuristics and biases view of probabilistic reasoning. It seems as though we need another view of students' conceptions of probability that is different from that offered by the traditional literature.

Results: A Framework for Understanding Probabilistic Reasoning

The main finding from this research is that students display a wide variety of ideas, some of which approach normative reasoning in probability, and others of which interfere with normative reasoning. In this study, instead of analyzing students' statements in an attempt to derive misconceptions, we will compare students' reasoning and normative probabilistic reasoning.

To make this comparison, we must first have an understanding of what we mean by normative reasoning. The version of normative reasoning used here is an idealized reasoning process used by someone with an understanding of elementary probability who is faced with a novel situation. The novel situation in this case is determining if the games of chance described previously are fair.

Such normative reasoning will first determine what is meant by "fair", which will not be addressed in this paper. After this, any of the several different reasoning processes that can be considered "normative" will have the characteristics that they will be based on (i) determining that the game is based on a non-determinable mechanism (i.e. understanding some aspects of randomness); (ii) determining the outcomes that score points for each team (i.e. understanding the outcome space); (iii) determining the probabilities of the outcomes that score a point for each team (i.e. understanding the probability distribution), and combining these probabilities with the outcome space to derive a theoretical expectation of fairness; and (iv) comparing the expected fairness of the games to the actual fairness after playing for some large number of points to determine if the theoretical expectations are accurate. In order to understand how students' reasoning differs from, and is similar to, normative reasoning, we will compare student reasoning to this idealized view of probabilistic reasoning.

Randomness

During the course of the study, every pair of students made reference to the fact that randomness was an important factor in analyzing the outcomes of coin flips. This reference typically came through students talking about the game being based on "luck" or "chance", and also by contrasting these games with games of skill. Additionally, students stated that the random process of coin flips would result in variation between trials (Figure 14), which is consistent with the normative view of randomness.

K: It won't be totally even
U: ... when you flip 2 coins you don't know what they are going to be. Sometimes you win, sometimes you lose.
X: The tides can change.

Figure 14: Randomness has variability

K (typed): I THINK THIS GAME IS FAIR BECAUSE IT'S JUST LUCK OR FATE WHAT IT LANDS ON!!!!!!!
O: Anything can happen, so that's why I think it's fair

Figure 15: Luck implies fairness

However, as illustrated in the case study of Q and T, three of the eight pairs of students stated that luck or randomness meant that nothing could be predicted about the games (see Figure 15). This result may be consistent with the "outcome approach" as described by Konold et al. (1993), as the students were replying to a question about a series of events as if the answer depended upon being able to predict any single event. These students did not accept that the Three-Coin game was unfair until they noticed the difference in the number of paths, and only then were able to create a new understanding that could explain the data.

The Outcome Space and Probability Distribution

In formal probability theory, determining the probability of a compound event is a multi-step process: one determines the relevant outcomes and, creating or using a probability distribution, one determines the likelihood of each of these outcomes. Combining these likelihoods determines the probability of specific events (such as Team A scoring a point). Using this process, one can clearly differentiate between the outcome space and the probability distribution, and these two entities are often introduced at different times in probability textbooks (e.g. Pitman, 1993). Although the students in this study did invoke ideas similar to the outcome space and the probability distribution, they often reasoned in a way that made it difficult to distinguish between the two.

Many students had trouble in differentiating the *individual* outcomes from the *set* of all outcomes that could score a point for a team. So, when asked the probability of a specific outcome occurring, many of the students stated that *each* of the outcomes that score a point for Team A were more probable, because they expected Team A to win. That is, the students had a difficult time understanding how to differentiate between the outcomes space and the probability of a complex event such as a team scoring a point.

Several of the students explicitly stated that order did not matter when differentiating between outcomes, and then discussed a probability distribution over combinations of outcomes (Figure 16). This behavior could easily be seen as representativeness, however, this behavior is based on not fully understanding how to properly enumerate and partition the outcome space. This is in contrast with students who don't expect "patterns" in data (Figure 17). The latter is closer to the traditional definition of representativeness, and is based upon applying a non-normative probability distribution to a normative enumeration of the outcome space.

M: ... see, like, there's three of HHT, and 3 of TTH, but only one of HHH and one of TTT, so it's hard to get HHH, it's easier to get HHT

Figure 16: Not differentiating outcomes based on order

D: Umm, I chose those I guess cause it's like too much of a pattern, like Tails tails tail uh...yeah, seems like less of a chance for it to go on that same one all three times, and like tails heads tails, like, in a pattern like that, and like heads tails heads

Figure 17: Representativeness based on a non-normative probability distribution

Finally, although most students did reference the outcome space in the Two-Penny game, only two pairs of students referenced the outcome space when making predictions in the Three-Coin Game, with only one of these pairs explicitly counting the outcomes. This may be because it was easy for students to understand and verbalize the outcome space for the Two-Penny game, whereas the outcome space of the Three-Coin game was more complex, making it harder to discuss. As a result, even though the outcome space was always on the screen for both games, students were rarely able to use the outcome space to reason normatively about the games.

The Validity of Evidence

It is in documenting people's beliefs about the validity of data that the misconceptions literature in probability is least controversial: it is well known that people, including trained statisticians, often fall prey to the law of small numbers (Tversky & Kahneman, 1982) (an interesting exception is found in the illusionary correlation literature, Chapman & Chapman, 1967). However, the data from our study presents a picture that is not as clear-cut as the existing literature would lead one to believe. Although many of the students did exhibit behavior consistent with the law of small numbers at some times, students also expected variability between different trials of a random process, and several students did not fall prey to the law of small numbers: in fact, they explicitly denied the relevance of data that was in contradiction to their theories, not accepting the data as relevant until they had created a scenario that could fit the data (note that this behavior is consistent with the science education literature (Chinn & Brewer, 1993)).

Although this paper is not the forum for an in-depth analysis of the interaction between students' prior expectations and the role they gave to evidence, it is worth mentioning that students can be roughly characterized as behaving in a manner that was either data-driven or theory-driven. Data-driven students were characterized either by their unwillingness to create a theory in the absence of data, or by their willingness to give up their theory after only a small number of points had been played (typically 10 points or less). Theory-driven students were characterized by an unwillingness to believe the data when it was in conflict with their theory. When theory-driven students first saw that the data was inconsistent with their theory, they explicitly denied the relevance of the data, and stated that the computer was cheating, or they just ignored the data and claimed the truth of their predictions (Figure 19). The students abandoned their contention that the Three-Coin game was fair only after they were able to construct an understanding of the importance of the outcome space.

L: In real life it would be fair...the computer is cheating

Figure 19: Theory-driven discussions of evidence

Summary

Middle school students do have productive ideas about probability, and these can be seen as a source of normative probabilistic reasoning. When discussing randomness, many of the students explicitly stated that variation is to be expected from trial to trial. However, some students also believe that a random event means that nothing at all can be predicted about future events. It is important to note, though, that this can be viewed as a normative understanding of randomness that has been over-extended. That is, a frequentist understanding of probability maintains that it is meaningless to apply probabilities to specific events, so these students are in many ways "correct" when applying this idea to short-run data. It is only when extending this idea to long-run data that this reasoning is non-normative.

When discussing the outcome space and the probability distribution, students exhibited many different ideas, making any generalizations difficult. We can say, however, that many students have difficulty in distinguishing between outcomes, and many students have difficulty in distinguishing between the outcome space and the probability distribution. And, once the relevance of the outcome space becomes apparent, many students display reasoning that is consistent with normative probability.

The research literature shows that many people fall prey to the law of small numbers, and this study is no exception. However, this study illustrates subtleties that have not been fully appreciated in the literature. In particular, the theory-driven subjects in this study did not believe data that was in conflict with a theory that they had proposed. This suggests that there is an interaction between people's expectations and the validity that they are willing to attribute to data. We posit that this has not been appreciated in the past due to the artificial nature of the tasks that subjects were given, whereas in our tasks students were engaged and felt ownership of their predictions and conclusions.

Finally, the understanding of student intuitions presented here shows that, even when students do behave in a manner consistent with representativeness, different students may be using different reasoning strategies. Using the framework in this paper, we can see that there are two very different causes behind similar behavior: in some cases, the student may not have a full appreciation for the outcome space, and may not be differentiating between outcomes based on order, whereas in other cases, the student may be using a non-uniform probability distribution, and may believe that certain outcomes are more likely than others.

Conclusions

Students have many ideas about probability, and these ideas are not adequately described by simply stating that students are using heuristics such as representativeness. Instead, students invoke a large number of intuitions about probability, and these intuitions can be seen to roughly correspond to randomness, the outcome space, probability distribution, and the role of data. By viewing students' probabilistic intuitions in this way we expect that, although many of the misconceptions found in the literature are adequate ways of describing the behavior of some students some of the time, students will exhibit great variation in behaviors, based on their understanding of these four related areas. In fact, this variation is exactly what is observed in this study as well as in the research literature. We feel that by viewing student ideas about probability as consisting of four interrelated sets of intuitions, we can come to a more thorough understanding of probabilistic reasoning.

Acknowledgments

We would like to thank Robert Constantakis for his help with PIE, as well as Barbara White and the members of the InterActiveMedia Study group for their comments. We would also like to thank Jesse Ragent for his help in recent designs of PIE. This research is funded in part by a grant from the UC Urban Community-School Collaborative.

References

Chapman, L.J. & Chapman, J.P. (1967). Genesis of Popular but Erroneous Diagnostic Observations. *Journal of Abnormal Psychology*, 72, 193-204.

Chinn, C. & Brewer, W. (1993). Students' Responses to Anomalous Data: A theoretical framework and implications for science instruction. *Review of Educational Research*, 63(1), 1-49.

Gigerenzer, G. (1996). On Narrow Norms and Vague Heuristics: A reply to Kahneman and Tversky. *Psychological Review*, 103(3), 592-596.

Hilton, D. J. (1996). Logic and Language in Causal Explanations. In D. Sperber, D. Premack and A. Premack, (Eds.), Causal Cognition: A Mulitidisciplinary Debate, pp. 495-525. Oxford: Clarendon Press.

Konold, C., Pollatsek, A., Well, A., Lohmeier, J., & Lipson, D. (1993). Inconsistencies in students' reasoning about probability. *Journal for Research in Mathematics education*, 24, 392-414.

Lajoie, S. Jacobs, V. & Lavigne, N. (1995). Empowering Students in the use of statistics. *Journal of Mathematical Behavior*, 14, 401-425.

Pitman, J. (1993). Probability. Springer-Verlag, New York.

Shaughnessy, J.M. (1992). Research in Probability and Statistics: Reflections and Directions. In D. Grouws (Ed), Handbook of Research on Mathematics Teaching and Learning, pp. 465-493. New York, NY: Macmillan

Tversky, A. & Kahneman, D. (1982). Judgment under uncertainty: Heuristics and biases. In Kahneman, Slovic, and Tversky (Eds.), Judgment under uncertainty: Heuristics and biases, pp. 3-22. New York, NY: Cambridge University Press.

White, B. (1993). Intermediate Causal Models: A missing link for successful science education? In R. Glaser (Ed.), Advances in Instructional Psychology, pp. 177-252, Hillsdale, NJ:Lawrence Erlbaum Associates.