Internet Enhancements for

"Are peasants risk-averse decision-makers?"

A. What economists mean by "risk aversion': decreasing marginal utility functions

Many economic anthropologists deal with risk and risk aversion in the same way as economists. Economists have attempted to capture the concept of "risk aversion" by formalizing the idea that as an individual gets more and more of something, she values each additional increment less and less. If we are giving you eggs, you might value the first two or three eggs quite highly (especially if you are planning breakfast), the 6th and 7th a bit less, and the 49th and 50th eggs hardly at all. Mathematically, economists describe such individuals as having concave utility functions-see Figure 1. In this approach, individuals select among alternative practices or options by computing the expected utility associated with each option, and then picking the choice with the higher expected utility. On Figure 1, the x-axis is a quantity of wealth: like cash, sacks of grain or M & M's. The curve allows us to convert from wealth to utility, which is measured on the y-axis and represents the quantity that individuals want to maximize.

To see how such a concave shape generates "risk-aversion," suppose a farmer knows his current cropping strategy will produce a yield of 30 sacks of wheat, but he is considering an alternative "green revolution" technology that produces 60 sacks of wheat half the time, and entirely fails the other half of the time (suppose it yields just enough to replace the seed). Which does he choose? Looking at Figure 1, we can compute the farmer's expected utility from option 1-i.e. 30 sacks for sure. Starting at a wealth of 30, one can trace up to the curve and over to the corresponding utility value, which is also about 30. Because this is a sure thing, with no chance of failure, the utility of 30 sacks equals the expected utility from this decision. Computing the EU (expected utility) for option 2, the green technology, is a bit more complex. The farmer determines the expected utility (EU) of this option in the following manner:

$$EU = U_1 \bullet PROB + U_2 \bullet (1 - PROB)$$

Where U1 is the utility the farmer gets from 60 sacks of wheat, U2 is the utility from zero sacks of wheat, and PROB is the chance of getting 60 sacks instead of zero sacks, which equals 0.50 in this case. To get U1 we trace up from 60 sacks of wheat on the x-axis to the curve, and over to the corresponding Utility (U), which is about 50. U2, the utility of zero sacks of wheat, is zero. Substituting into the above equation gives an expected utility of 25, which is less than option 1's utility of 30, so our farmer prefers option 1-he prefers the sure thing (i.e. his traditional practices) over the variable outcome (the new technology).

This farmer is "risk averse" because of the concave shape of his utility curve. In contrast, if his utility curve were a straight line, he would be indifferent between the two options described above because each would produce the same expected yield of 30 sacks-this is "risk neutral," A risk-prone farmer would have a convex utility curve (the curve would accelerate upward, instead of decelerating downward), so he would always prefer the risky option that yields 60 sacks half the time.

B. Satisficing or Safety-first: minimizing the chances of falling below a subsistence threshold

Another way researchers use the concept of “risk aversion” is to suggest or hypothesize that farmers, very small birds, horticulturalists and/or foragers make decisions (i.e. select among alternative actions) in order to minimize their chances of falling below some subsistence minimum—which may be culturally or biologically defined, depending on the researcher (Winterhalder 1990; Winterhalder *et. al.* 1999; Real & Caraco 1986; Johnson 1971; Schultz 1964: 31). With some plausible assumptions, we can formalize and clarify this often-intuitive approach. On Figure 2, curves 1 and 2 represent the normal distributions of crop yields for two kinds of wheat (say winter and spring wheat). The symbols μ_1 and μ_2 stand for the average yields of these two varieties, while *R* represents the minimum amount of wheat harvest needed to avoid catastrophic consequences like starvation, the death of a child, the loss of land, etc. Rather than maximizing their average yield, which would cause them to sow crop 2, risk-averse farmers may prefer crop 1 because, although its average yield is less than crop 2 (μ_1 < $(_2)$), the probability that one falls below the catastrophic threshold, *R* is less for crop 1 than crop 2. To see this graphically, look at the area under each curve to the left side of R—that area equals the probability of falling below R, and suffering the catastrophic consequences. This conceptualization is common in both behavioral ecology in the analysis of risk-sensitive foraging (Stephens & Krebs 1986) and in economics for the analysis of capital asset markets and portfolio composition (Nicholson 1995). See Roumasset (1976) for a discussion directly related to farming risk.

This conceptualization captures what many researchers mean when they say peasant farmers are "risk averse." That is, they mean peasant farmers cannot adopt a new type of seed, planting practice or fertilizer (for example), which they believe may raise their average yields, because new practices have (either objectively or subjectively) substantial variation in outcomes. Any new technique or practice, even if its long-term variance is small, will often have a high variance at start-up, because such novelties often require experience and/or experimentation. During this period, the farmer must refine his skills and understanding of the new practice in adapting it to the particulars of the local situation. From this perspective, given their economically precarious situation, many peasants might, quite sensibly, prefer tried-and-true practices (which consistently keep them above their minimum threshold) to new practices with greater variation. This all makes sense, but is it a good model for peasant behavior?

C. The Titration Experiment

We detail the Titration methodology using Henrich's experiment—McElreath's procedure was similar unless otherwise indicated. Henrich performed the *Titration Experiment* with 26 Mapuche farmers and 25 Huinca townspeople. For standardization purposes, Henrich used a Spanish-language version of the form shown in Figure 3. In all these experiments, Henrich followed the same basic script. Although Henrich used a written version of the game, it did not matter if subjects could read Spanish. Henrich explained the entire game verbally and used the sheets only to remind subjects of the numerical values involved—all participants seemed to understand the numbers. First, if he had not previously worked with the person, he was introduced to them by a mutual friend. Henrich explained that he was administering an "economics experiment." Because he had been working with, interviewing and/or interacting with all of the Mapuche participants for several months prior to administering this experiment, introductions were mostly unnecessary, and few people asked much about the experiment. For many of the Huinca participants, however, this brief introduction was necessary because Henrich did not know them, or had only a passing familiarity with them. Second, Henrich explained the game using the form (Figure 3) as a visual aid. Henrich emphasized that all the rounds were independent. He explained that participants would first make all their choices in the three rounds, and then at the end, they would "play" any risky choices (i.e. flip the coin) and he would pay them the total amount earned in all three rounds. Henrich also showed the participants an envelope full of cash and coins, and assured them that all winnings would be paid in full right after the game. Henrich explained round 1 as follows (this is translated nearly verbatim):

> In round 1 you have to choose between option A and option B [pointing to each]. Option A is 1000 *pesos* for sure. I simply pay you 1000 *pesos* in option A. In option B, you have a 50% chance at 2000 *pesos*, but you may lose and receive nothing. If you pick option B, you will flip a coin. If the coin comes up "heads" you get 2000 *pesos,* but if tails comes up you get nothing. If you pick option A, you get 1000 *pesos*—no coin flipping. If you pick option B, you flip the coin—'heads" you get 2000 *pesos* and "tails" you get zero *pesos*.

Then, he asked if the participant had any questions. If they did not, Henrich would then ask them three test questions: (1) "if you pick option A, how many *pesos* do you get?" (2) "if you pick option B and tails comes up, how many *pesos* do you receive?" and (3) "if you pick option B and heads comes up, how many *pesos* do you receive?" If the participant answered all three questions correctly, he would then ask them which option they prefer, A or B. If they did not answer correctly, he re-explained the game, and we would try again. Most people understood the game fairly easily and got all the questions right the first time. A few older Mapuche never got the correct answers, and Henrich eventually had to give up, and not use them in the experiment.

If the participant picked the risky bet (option B) in the first round, Henrich would "sweeten" option A in the next round by increasing it from the 1000 *pesos* to 1500—which Henrich would write down in the blank next to round 2 (see Figure 3). If the participant picked the safe bet (option A) in round 1, Henrich would "sour" option A and put "500" in the blank space in round 2. The idea being to either sweeten or sour the safe bet until the participant switches from the risky to the safe bet, or from the safe to the risky bet. Round 2 was administered much like round 1. If the participant picked the risky bet on round 2, Henrich would increase the value of the safe bet by 300 *pesos* in round 3, by writing down either "1800" or "800," depending on the participants previously choices. If the participant picked the safe bet on round 2, Henrich would decrease the value of option A to either 1300 or 300 *pesos*—again depending on their previous choices. If subjects could have foreseen what Henrich was doing with the bets, they might have calculated that the best strategy is to pick BAA. However, it's very difficult to see how subjects could have deduced this *apriori*, even if they had talked among themselves. It was also clear from the subject's actual behavior that no one figured this out.

Indifference points were recorded as the amount halfway between the sure bet in round 3 and the nearest known decision point. For example, if a participant picked "A" in round 1, round 2 becomes a choice between 500 *pesos* for sure and 2000 with a 50% chance. If the participant then picks "B," round 3 becomes a choice between 800 *pesos* for sure or 2000 *pesos* with a 50% chance. If the participant then picks "B" again, the indifference point is recorded as 900 *pesos* (he picked "A" when it was 1000 and "B" when it was 800 *pesos*). Also note that participants who picked BBB were recorded as 1900, and participants who picked AAA were recorded as 150 (150 is halfway between 300, where AAA subjects chose the safe bet, and zero). After round 3 was complete, the participant flipped the coin for any risky bets and Henrich paid them the total amount owed. After this, he interviewed each participant about why they made particular choices.

D. The Variance Experiment

The basic structure of the *Variance Experiment* is very similar to the *Titration Experiment*. The goal was to gather data on how changing the variation in outcomes affects people's economic decisions, when the expected value of the options remains the same. Figure 4 depicts one of the two forms used to play the game, which Henrich administered to both Mapuche (41 participants) and UCLA undergraduates (20 participants). The expected value (average return) of both options in all four rounds is the same (1000 *pesos* for Mapuche and $15 for UCLA undergraduates), except in round 4b, where the chances of winning were unknown. What does change from round to round is the probability of winning and the amount won in option B (i.e. the variance), not the expected return.

Henrich administered round 1 of this game in the same way as round 1 of the *Titration Experiment*, except that risky bets were played right away (rather than waiting until the end) and players were paid after each round. For the 50% bet, Henrich used the same coin-flipping method as in the *Titration Experiment*, but for the other gambles participants had to pick a white card out of a hat. For the 20% and 80% gambles, Henrich used the same five cards, four with "X's" and one with a "Z'. In the 20% case, if the participant picked the "Z" out of the hat, he received 5000 *pesos* (or $75 for UCLA students). If he chose an "X," he got zero *pesos*. In the 80% case, if the player chose one of the four "Xs," he received the 1,250 *pesos* (or $18.75), but if he chose the "Z," he received nothing. For the 5% gamble, 19 blank cards were counted out, one "X card" was mixed in, and the whole pile was put into the hat. If the player chose a blank card he received no money. If he picked the "X," he got 20,000 *pesos* (or $300 at UCLA).

In some versions of the game, instead of choosing between 1000 *pesos* for sure or 20,000 *pesos* with a 5% chance, round 4 provided a choice between 1000 or 5000 *pesos* with an *unknown* chance. For this game, Henrich removed small stack of cards from his knapsack and told the participant that the stack contains some blank cards and some "Xs." Henrich would explain that: "if you pick an "X" you get 5,000 *pesos*; if you pick a "blank," you receive zero *pesos*; but, you do not know how many "X's there are nor how many "blanks" there are in this stack." Players would then pick either option A or B. This alternative round 4 was not performed at UCLA.

As a method of controlling for the order of presentation of the different gambles, Henrich also performed a number of games among the Mapuche using another form in which the first three gambles were presented in order 80%, 50%, 20%, instead of the 50%, 20%, 80% shown in Figure 4. Round 4 or 4b was always presented last. Henrich administered only the version shown in Figure 4 to UCLA undergraduates.

In McElreath’s work among the Sangu of Tanzania, he used the same procedure, except for two things. First, he did not use forms; he simply explained the game verbally, though he always presented the gambles in the order shown in Figure 4. Second, he used two cards to generate probabilities in the 50/50 gambles rather than flipping a coin. McElreath also matched Henrich’s stake size by using the same calculation method described above.

E. Analysis of empirical concerns

For the remainder of this section, we address five questions about this experimental data: 1) Can the observed difference in risk preferences between the groups be explained by the different proportions of males and females found in the samples? 2) Does the order of presentation of the different gambles matter in the Variance Experiment? 3) Do people look further down the form and adjust their choices according to what they see? 4) Do people change their answers in the later rounds because of wins or losses in the earlier rounds (i.e. are the rounds independent)? And, 5) Did people understand the games? In the next section, we focus on the implications of these results for understanding peasant behavior in terms of the different ways researchers have used “risk aversion’.

1) Can the difference in risk preferences between the groups be explained by the different proportions of males and females found in the samples? No. Because Henrich’s research among the Mapuche focused primarily on agricultural practices, Henrich administered the V*ariance Experiment* to mostly heads-of-household. There were 36 male heads-of-household and only five female heads. These five females performed quite similarly to the males, but of course, there are only five. Compiling all the choices from all four possible bets, Mapuche males (36 people and 144 choices) picked the risky bet 73% of the time, while Mapuche females (5 people and 20 choices) took the risky bet 90% of the time. The addition of the females to the male sample moves the average from 73% to 75%. Males and females are not significantly different, although females tended to be more *risk-prone* than males.

Among UCLA students, 14 males and 6 females revealed quite similar behavior in the *Variance Game*. Table 1 summarizes the results for males and females across the different gambles. We do this for the UCLA sample and not the Mapuche sample because, unlike Mapuche, the behavior of UCLA students varied greatly with the amount of risk or variance in a gamble. Table 1 shows that UCLA males and females, like Mapuche males and females, behave quite similarly. In both samples the females make the group average slightly more risk-prone. This means that if the difference between the groups was driven by the differences in female/male ratio, then the group with more females should be more risk-prone. In contrast, the UCLA sample, which has a higher female/male ratio than the Mapuche sample, is much more risk-averse. Consequently, if the difference in sex ratio between the samples affects the results at all, it makes the two groups seem more similar. Furthermore, for all of our groups, multivariate logistic regressions using sex, age, head-of-household, and various measures of wealth/income, consistently show no significant effect of a subjects' sex.

The within-group analysis of the Mapuche and Huinca titration data further indicates that the small differences in the sex ratio used in the two samples does not seriously influence the large differences observed between the Mapuche and the Huinca. Table 2 summarizes the breakdown. The Mapuche *Titration Experiment* used 22 males and 4 females. There was no significant difference between them, though the females tended to be more risk-prone than males (again). Similarly, the Huinca sample consisted of 18 males and 7 females, and also shows no significant difference—although here females tend to be more risk-averse than the Huinca males. Removing the females entirely from both samples still reveals a substantial difference in the mean indifference points for each group. Furthermore, the multivariate regressions in Tables 1 and 2 from the main text show that sex is not an important predictor of risk preference. The point here is not to argue that the sex of the participant is irrelevant, but only to claim that the differences in the sex ratios used in the experiments cannot explain the observed differences between the groups, and are too weak to detect.

Experimental work done by Binswanger (1980) using similar monetary gambles among peasant farmers in India further confirms that sex is not an important explanatory variable. In this work, participants have to pick a gamble to "play" from a series of possibilities that differ in their expected returns and variances. Essentially, participants had to trade-off higher variance outcomes against gambles with higher expected returns. Participants were paid in real money, and the stakes ranged from one-seventh of a day's wage to 14 days wage. Using this approach, Binswanger found that males and females have quite similar risk preferences, once educational level is controlled for. Using university subjects, Holt and Laurie (2000) also show that sex is not an important predictor variable when non-trivial stakes are on the line.

The Sangu show some interesting sex differences. Among Sangu farmers, females were significantly more risk-averse than male farmers. Yet among Sangu herders, both males and females remain risk-prone. If we extract only the male farmers from the Sangu data, and compare them against the Mapuche and UCLA data (where we have already shown that females do not significantly alter the results), we get Figure 5. This shows that male Sangu farmers are slightly more risk prone than the overall Sangu result, except in the low variance bet (80% chance of winning), where they drop down to the nearly risk-neutral position of UCLA students. In general, Figure 5 here is quite similar to Figure 2 in the main text.

(2) Does the order of presentation of the different gambles matter in the Variance Experiment? No. In the Mapuche *Variance Experiment,* Henrich performed two versions of the game. With 28 participants Henrich administered a version with the following sequence of winning chances 50%, 20%, 80% on the first three rounds. Of these 28 games, 17 ended with the unknown gamble for 5,000 *pesos* ("4b" in Figure 4), while the rest ended with the 20,000 gamble ("4" in Figure 4). With the other 12 participants, Henrich did a different version in which the order of the first three gambles was 80%, 50%, 20%. All 12 of these ended with the unknown bet. Table 3 shows the frequencies of risky bets and *p*-values for comparisons of the two game-versions. All of the means appear quite close and none approach typical significance levels—although a power analysis indicates that with sample sizes of only 12 and 28 statistical analysis won't become weakly significant ($p < 0.10$) until the difference between the two samples exceeds approximately 0.21. That is, the order of gambles may matter, but it's effect is less than 0.21, which is small relative to the cultural differences observed between groups (which is about 0.6). Our fixed-effects regressions also show no linear effect of round, except for the UCLA sample, in which order might be important.

(3) Do people look further down the form and adjust their choices according to what they see? Not too much. Both the *Titration Experiment* and *Variance Experiment* (50-20-80 version) began with the 50/50 bet for 2000 *pesos* or a certain 1000 *pesos*. In the *Variance Game*, participants may have glanced down the form, saw the other gambles, and altered their behavior with some inferences about what was to come. In contrast, Titration participants might have also looked down the form, but all they would have seen were incomplete gambles with blank lines (see the methodology section). Consequently, if this "glancing down" is important, we should expect the choices made in the first 50/50 bet in the *Titration* and *Variance Games* to be different. Instead, the frequency of risky bets is quite similar in the two games: *Variance Game* participants selected the risky option 73% of time ($n = 41$), while titrators picked it 80% of the time ($n = 26$; $p = 0.38$). In the Sangu games, there were no forms.

(4) Do people change their answers in the later rounds because of wins or losses in the earlier rounds? (i.e. are the rounds non-independent?) No. Table 4 addresses this question with a series of conditional probabilities calculated from the *Variance Game* data for both the Mapuche and UCLA experiments. Each frequency represents the fraction of times a participant took the risky bet given the following three (all-encompassing) conditions: (1) they had picked the safe bet on the previous round; (2) they had picked the risky bet and won the gamble on the previous round; and (3) they had picked the risky gamble and lost on the previous round. For obvious reasons, the first round was not included in this analysis.

Looking at the Mapuche and Sangu, it's clear that what happened on the previous round did not affect their decisions on the current round. The chances of picking the risky bet are quite similar in all three conditionals. If an individual lost the previous gamble, his chance of picking the risky option is 75%, while if he won his chance is 78%. At UCLA, having won or lost the previous gamble also does not significantly influence one's chance of taking the risky bet, although having taken the safe bet in the previous round doubles one's chances of taking the risky bet on the current round ($p = 0.04$). UCLA choices are independent of winning and losing on the previous round, but students seem to like to "mix things up" in choosing A or B (which could be interpreted as a kind of risk-aversion).

(5) Did people understand the games? Both hypothetical test questions and interviews indicate that Mapuche, Huinca, Sangu and UCLA students all understood the games. As we mentioned earlier, everyone had to answer the three test questions before playing each round. Most people got these correct right away. The few people who were never able to answer correctly were not permitted to participate. In the post-game interviews, for the both the Mapuche and Huinca, Henrich asked questions such as "in which round would you have the best chance of winning if you picked option B" or "in round 2, are you more likely to win or lose if you pick option B." Participants, for example, definitely knew that the chances of picking one card out of five were less than the chances of getting "heads" with the coin. People's answers were mostly correct, and indicated that everyone had a qualitatively accurate understanding of the game.

Section F.

Table 5 shows the fixed-effects logistic regressions using all 4 rounds of Variance Game data. Fixed effects regression uses dummy variables to control for differences among individuals, allowing us to analyze the effects of some treatment to which individuals were exposed. In this case, the treatments were round order and the variance of the bets. We see that neither round nor gamble variance have any effect on Sangu or Mapuche behavior. However, among UCLA students both gamble variance and round are important predictors. Each additional round, UCLA subjects were about 50% less likely to take the risky bet ($p = 0.02$), controlling for gamble variance. Gamble variance itself went from 0.8 (low variance) to 0.05 (high variance). On average, each 5% reduction in variance (moving towards a non-variable bet of 1.00) increased the odds of taking the risky bet by about 24% ($p < 0.01$). UCLA subjects are clearly sensitive to risk variance in these gambles in a way in which the other samples are not.

Figure 1. Decreasing marginal utility curve (risk aversion)

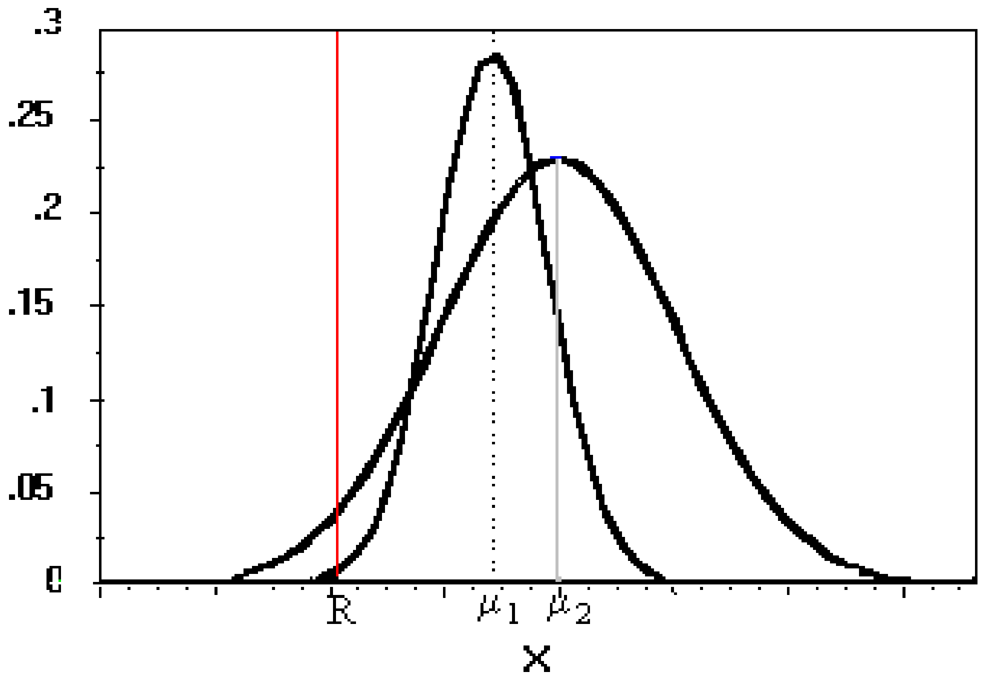

Figure 2. Normally-distributed crop yields for the satisficing or safety-first risk model. *R* is the minimum threshold. μ_1 and μ_2 are the mean yields of crop 1 and crop 2.

Figure 3. The Titration Experiment

The Preference Game

The are 3 rounds in this game and individuals play alone. In each round, you have to pick between 2 possibilities.

Do you prefer:

	A		B
1.	1000 *pesos* (100%)	or	2000 *pesos* (50% chance)
2.	_____*pesos* (100%)	or	2000 *pesos* (50% chance)
3.	_____*pesos* (100%)	or	2000 *pesos* (50% chance)

***The money is real and you will be paid.**

**There are no correct or incorrect answers. Pick whatever you prefer.

Figure 4. The Variance Experiment

The Risk Game

The are 4 rounds in this game and individuals play alone. In each round, you have to pick between 2 possibilities.

Do you prefer:

	A		B
1.	1000 *pesos* (100%)	or	2000 *pesos* (50% chance)
2.	1000 *pesos* (100%)	or	5000 *pesos* (20% chance)
3.	1000 *pesos* (100%)	or	1250 *pesos* (80% chance)
4.	1000 *pesos* (100%)	or	20,000 *pesos* (5% chance)
4b.	1000 *pesos* (100%)	or	20,000 *pesos* (unknown chance)

***The money is real and you will be paid.**

**There are no correct or incorrect answers. Pick whatever you prefer.

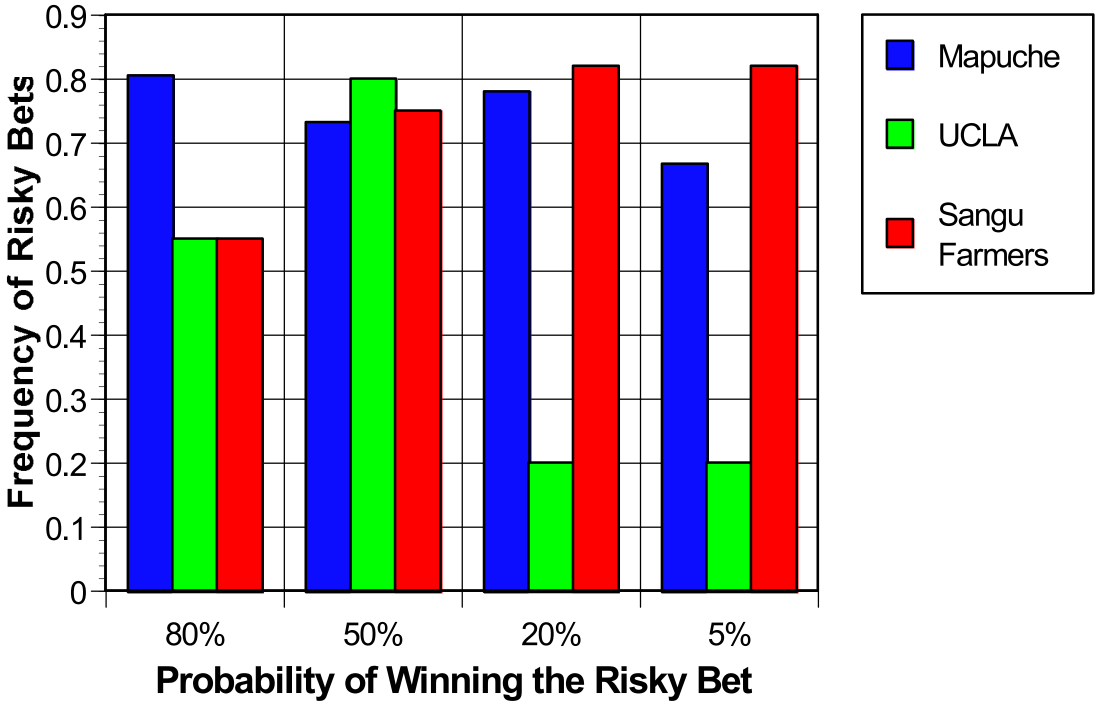

Figure 5. Variance results for Mapuche, UCLA and Sangu male farmers. The Sangu sample of male farmers has n = 11 for 80%, 20% and 5% gambles, and n = 20 for 50% gamble. The Mapuche and UCLA data is the same as in Figure 4.

Table 1. Frequency of Risky Gambles for UCLA, split by sex

UCLA	50 Gamble	20 Gamble	80 Gamble	5% Gamble	Avg. of Gambles
Males	0.79	0.14	0.50	0.21	0.41
Females	0.83	0.33	0.67	0.17	0.50

Table 2. Indifference Points for Mapuche and Huinca males and females

Group/Sex	**No. of Participants**	**Indifference Point***
Mapuche females	4	1525
Mapuche males	22	1377
Huinca females	7	721
Huinca males	18	816

**p*-values for both for male-female, within-group, comparisons are about 0.57

Table 3. Variance Game Data from 2 different bet orderings

Variance Bet	Risky Bet Freq. 50-20-80 version (*n*)	Risky Bet Freq. 80-50-20 version (*n*)	Comparison *p*-values	Freq. risky bet overall
80%	0.82 (28)	0.83 (12)	0.93	0.82
50%	0.79 (28)	0.67 (12)	0.44	0.75
20%	0.71 (28)	0.92 (12)	0.17	0.77
Unknown	0.88 (17)	0.67 (12)	0.17	0.79
Total*	0.79 (101)	0.77 (48)	0.73	0.78

*this calculation assumes each *round* is independent. Analysis presented in the next two sections justify this assumption.

Table 4. Conditional Frequencies of Risky Bets

Freq. risky bet given previous:	**Mapuche**	**UCLA**
Picked safe bet on previous	21/28 (75%)	13/29 (45%)
Picked risky bet and won	36/46 (78%)	4/20 (20%)
Picked risky bet and lost	36/48 (75%)	2/11 (18%)

Table 5. Fixed-effects regression using Variance Game data for the Mapuche, Sangu and UCLA students.

Sangu	**β**	**S.E.**	**Wald**	***p*-value**	**Exp(B)**
ROUND	0.1302	0.3262	0.1594	0.6897	1.1391
VARIANCE	0.4361	1.2773	0.1166	0.7328	1.5467
Constant	0.6064	1.6818	0.13	0.7184	

Mapuche	**β**	**S.E.**	**Wald**	***p*-value**	**Exp(B)**
ROUND	0.107	0.2795	0.1467	0.7017	1.113
VARIANCE	0.3852	1.0154	0.1439	0.7044	1.4699
Constant	10.805	156.1129	0.0048	0.9448	

UCLA	**β**	**S.E.**	**Wald**	***p*-value**	**Exp(B)**
ROUND	-0.7217	0.3064	5.5482	0.0185	0.4859
VARIANCE	3.206	1.0994	8.5036	0.0035	24.681
Constant	0.4535	1.4714	0.095	0.7579	