Signaling strength? An analysis of decision making in *The Weakest Link*

Marco A. Haan · Bart Los · Yohanes E. Riyanto

Abstract   We analyze contestants' behavior in the game show "The Weakest Link". We focus on banking decisions, where a contestant chooses to secure an amount of money for the eventual winner, or to risk it on a general knowledge question. We find that contestants do not use the banking strategy that maximizes total expected prize money. Average earnings could be at least 17% higher. Our results suggest that contestants are not overconfident, but do try to convince other contestants that their ability is higher than it really is, in order to increase chances of winning the prize. We argue that this mechanism may also be applicable to other situations that are of economic interest.

Keywords   Game show · Field experiments · Signaling

1 Introduction

This article studies decision making in the context of a TV game show, *The Weakest Link*. In this show, contestants repeatedly have to choose whether to secure ('bank') an amount of money for the eventual winner, or to risk it on a general knowledge

M. A. Haan (✉) · B. Los
Faculty of Economics and Business, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
e-mail: m.a.haan@rug.nl

B. Los
e-mail: b.los@rug.nl

Y. E. Riyanto
Division of Economics, School of Humanities and Social Sciences, Nanyang Technological University, 14 Nanyang Drive, HSS #04-83, Singapore 637332, Singapore
e-mail: yeriyanto@ntu.edu.sg

question. We test whether players use the optimal banking strategy. To do so, we study the banking behavior of 693 contestants. First, we try to find some general patterns in that behavior. Second, we test whether contestants use the strategy that maximizes total prize money. Third, we investigate how contestants' behavior systematically deviates from that strategy. Our results suggest that, different from what an increasing body of literature in psychology and economics suggests,[1] contestants do not suffer from overconfidence. Rather, they merely try to convince others that they have a high ability.

TV game shows form an ideal laboratory to study decision-making. The rules of the game are well-defined, and the stakes are often substantial. A number of papers have analyzed the behavior of contestants on game shows, for instance Berk et al. (1996), Gertner (1993), Metrick (1995), and Beetsma and Schotman (2001). The show *The Weakest Link* that we study has also attracted attention. Both Levitt (2004) and Antonovics et al. (2005) use data from the voting stage of this game to test theories of discrimination. Février and Linnemer (2006) study equilibrium selection in the final stage of the game.

The remainder of this article is structured as follows. In Sect. 2, we give a full description of *The Weakest Link*: the set-up, the decisions contestants have to make, and other relevant issues. Section 3 gives some descriptive statistics of the episodes we study. In Sect. 4, we investigate contestants' banking behavior. We show that, on aggregate, the contestants' propensity to bank follows some simple and intuitive comparative statics: contestants are more likely to bank when questions are more difficult to answer, when the amount at stake is higher, and when they have less general knowledge.

In Sect. 5, we test whether contestants use the banking strategy that maximizes total prize money. We show that an extremely simple decision rule could already have done better on average. To get a handle on exactly how much better contestants could have done by using the optimal strategy, we use a computer algorithm in Sect. 6. We find that, first, contestants could have increased their earnings by at least 17% and, second, that the intuitive comparative statics we found in Sect. 4 are *not* a part of the optimal strategy.

Using our algorithm, we set out to study in Sect. 7 to what extent contestants' behavior differs from the optimal strategy. We find that during the first seven rounds of the show, in an overwhelming majority of cases, contestants make 'wrong' decisions by banking too late, rather than too early. In round 8, this picture changes and we observe an overwhelming majority of cases where contestants bank too early rather than too late. Yet, as we will argue, in round 8, contestants have no reason to try to convince the others of their general knowledge, whereas in the first seven rounds they do have such an incentive. We thus conclude that, on average, contestants of *The Weakest Link* bank later than they should in an attempt to try to convince the other players that they are more knowledgeable than they really are. In other words, contestants try to signal to the other players that they are confident about their general knowledge. Yet, contestants do not suffer from overconfidence: they do realize that they are not as good as they suggest to be. We argue that this mechanism may also be applicable to

---

[1] See e.g., Camerer and Lovallo (1999) or Bénabou and Tirole (2002) and the references therein.

other situations that are of economic interest. In any situation where an individual is evaluated based on his perceived quality, he may be inclined to take decisions that are too risky, in the sense that they are the optimal choices of an individual of higher quality. We give some examples in Sect. 8, which also concludes.

2 The game

The Weakest Link is a daily game show that has been broadcast by the British public television network BBC since August 2000. It has rapidly become one of the most popular game shows in the UK. Local versions of The Weakest Link are (or were) shown in dozens of countries, most of them with virtually the same set of rules as the UK version. Part of the success must be attributed to the unorthodox attitude of the quizmaster. Instead of showing empathy and support, (s)he berates the contestants throughout the show by giving sarcastic comments on their supposedly non-existent intelligence and personal traits such as an odd hairstyle.

The show starts with nine contestants of whom only one will take home any prize money.[2] In every round, the contestants receive trivia questions. The players stand in a semi-circle around the quizmaster. Each contestant receives one question at a time, moving in clockwise order. A correct answer yields a 'link' in a 'prize chain', beginning at zero and climbing to £1,000 in nine increments, as shown in Fig. 1. Whenever a player gives a wrong answer, the chain is broken and all money on the chain is lost. Questions are *not* more difficult to answer when the amount of money on the chain is higher.

However, it is not necessary to reach £1,000 on the chain to secure winnings and add to the prize money. Just before a contestant receives a question, he has the option to 'bank' the money on the chain by saying the word 'bank'. The amount on the chain at that time is then added to the prize money, and the chain drops back to zero. After banking, the contestant still has to answer a question. Of course, if the amount of money on the chain is zero, there effectively is no banking decision to be made.

A round ends when time has run out. Any money on the chain that has not been banked before the end of the round is lost. The first round lasts 3 min. Each subsequent round lasts 10 s shorter than the preceding one up to round 7. In round 8, the two remaining contestants have 90 s. A round may end before time is up when a total of £1,000 is banked.[3]

As an example, consider the following chain of events in the first round. The first contestant is posed a question and gives the correct answer. If contestant 2 banks, an amount of £20 is secured. The chain moves back to 0. Now contestant 2 receives a question. She answers correctly as well. Contestant 3 does not bank and gives the correct answer. The chain has then moved to £50. Contestant 4 does not bank and gives the wrong answer. The chain drops back to zero, and the money on the chain is lost. Contestants 5 and 6 do not bank, but do give the correct answer, moving the

---

[2] Note that the description in this section applies to the BBC shows we studied. Shows in some other countries (including the US and the Netherlands) only have eight contestants.
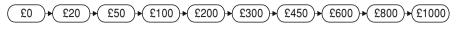
[3] If the amount banked in one round exceeds £1,000, the excess will not be added to the prize money.

Fig. 1 The prize chain

chain back to £50. Contestant 7 banks, adding the £50 to the secured money, which now stands at £70. At the end of a round, the amount of money secured in that round is added to the total prize money of the show.

After each round, except round 8, the team gets rid of one player through a voting procedure. The person with the largest number of votes receives one more sneer from the host and is then sent home with the infamous words "You are the weakest link. Goodbye!" In case of a tie, the person who answered the most questions correctly in that particular round—the "strongest link"—casts the deciding vote, which may be different from her original choice. The idea of the voting procedure is that contestants should vote for the player they think is the worst among those remaining. This player would be least likely to contribute much to the prize money.

Evidently, the voting procedure may induce strategic behavior. In round 7, for example, with only three players left, there is a strong incentive for the two weaker players to vote off the strongest link, giving them a greater chance of winning the final. Yet, the trade-off is that prize money in the final round may be higher if the strongest player is still in the game. To try to weaken this incentive for strategic voting, the amount banked by the two players is trebled in round 8.

Each episode ends with the final, in which the two surviving contestants play against each other head to head. The host alternately asks each contestant five questions. In case of a tie after both players have answered five questions, more questions are posed to determine who wins. The winner takes home the prize money accrued during the eight rounds. The loser leaves with nothing.

3 Descriptive statistics

For this article, we watched 77 episodes of *The Weakest Link*, originally broadcast on weekdays by the BBC in the period December 2000 through May 2001,[4] relatively early in the show's history. For every episode, we recorded all voting decisions in every round and for every question posed, the identity of the player answering, the amount of time left on the clock just before a question was posed, whether the question was answered correctly, and the banking decision. Table 1 gives some descriptive statistics of our data set.

The probability of a correct answer is defined as the total number of correct answers divided by the total number of questions. The propensity to bank is defined as the number of actual banks divided by the number of times contestants could have banked (i.e., the number of correct answers). Thus, on average, contestants bank in 29% of all pos-

---

[4] Note that this period has far more than 77 weekdays. Some episodes broadcasted towards the end of our sample period were reruns of earlier broadcasts. On public holidays, celebrity editions were often aired, which we did not include in our data set. Also, the broadcasting schedule was sometimes interrupted to show sports events instead.

Table 1 Some descriptive statistics

Episodes	77
Questions per episode	164.74
Correct answers per episode	96.46
Number of banks per episode	28.34
Average prize money per episode	£2,339
Minimum	£1,080
Maximum	£3,820
Average time per question	6.92 s
Probability of answering correctly	0.586
Propensity to bank	0.294
Average money earned per bank	£75.40
Average money earned per correct answer	£22.15
Average number of correct answers at bank	2.14

sible cases. The "average number of correct answers at a bank" equals the average number of steps up the chain when a bank occurs. Thus, on average, there is a chain of 2.14 correct answers when a contestant banks.

Note that the average time of 6.92 s per question is very short: during this time span contestants not only have to listen to and answer their question, but also have to make a banking decision. Moreover, they need to keep track of the performance of their opponents. All of these, plus the fear of a ruthless putdown by the host of the show in case of failure, imply that contestants are subject to considerable stress. This will have an effect on their performance, both in answering the questions, and in making banking and voting decisions.

Table 2 gives some additional statistics, broken up to individual rounds. As the show progresses, it becomes more difficult to answer a question correctly. This may be because the questions themselves become more difficult,[5] or because players are increasingly affected by stress and fatigue. For our analysis, this is immaterial. In round 1, contestants managed to bank the maximum amount of £1000 in 17 episodes. In round 2, this happened in three episodes.[6] In later rounds, there are no such cases.

This article focuses on the banking decision, which is a non-trivial optimization problem. A strategy of always banking late (i.e., if the amount on the chain is high) has the potential of yielding high winnings, but also implies the risk of losing a substantial amount of money if a wrong answer is given. A strategy of always banking early implies that only little money will be earned per bank. Also, there is an option value to

---

[5] Note that this does *not* contradict the earlier statement that, *within a given round*, questions are not more difficult when the amount of money at stake is higher.

[6] During our sample period, there seems to have been a change of policy with respect to the difficulty of the questions in the first round. During the first 30 episodes in our sample, there has not been a single case in which the £1000 target was reached in round 1. During the other 47 episodes, this occurred 17 times. Such systematic differences are not found for later rounds.

Table 2 Descriptive statistics per round

Round	1	2	3	4	5	6	7	8
Duration of round (s)	180	170	160	150	140	130	120	90
Probability of answering correctly	0.764	0.687	0.606	0.570	0.516	0.483	0.434	0.446
Average prize money	638	418	311	246	159	141	123	101
Minimum	90	40	20	50	0	0	0	0
Maximum	1000	1000	900	850	450	400	460	310
Occurrences of 1000 banked	17	3	0	0	0	0	0	0

banking late: doing so not only increases the amount of money currently on the chain, but also allows contestants to reach even higher amounts on that chain.

Ultimately, the expected earnings of a single contestant in *The Weakest Link* depend on the knowledge of all contestants, and on their banking and voting decisions. We assume that the goal of a contestant on the show is to maximize these expected earnings.[7] It is conceivable that voting decisions are based on banking decisions of the other players. Contestants may be more likely to vote for someone whom they feel to use a non-optimal banking strategy. If this is the case, then banking decisions will be chosen not only to maximize expected earnings, but also to influence voting decisions of others, and hence to stay in the game. As our working hypothesis, we assume that this is not the case. Thus, we assume that banking decisions are taken solely with the aim to maximize expected earnings. Effectively, we assume that contestants take their banking decisions as a group, since their objectives are perfectly aligned. Given that the ultimate winner will take home all of the winnings and everyone has a positive probability of being that winner, it is in each player's interest that total prize money is as high as possible.

4 An analysis of contestants' behavior

We now take a closer look at the contestants' banking behavior. We show that, at least on aggregate, candidates behave according to the following intuitively appealing rules:

Result 1. *Contestants in our sample of BBC episodes of* The Weakest Link *are more likely to bank when*

(a) *the amount of money at stake is higher*,
(b) *questions are more difficult to answer*,
(c) *a player has less general knowledge*.

Note that the first result makes intuitive sense: if the amount of money at stake is higher, the amount of money lost if a wrong answer is given, is higher as well. Also,

---

[7] Admittedly, other objectives may play a role as well, such as the desire to appear sympathetic or knowledgeable on national television, to outsmart the host, etc. Yet, we assume that these objectives are only minor compared to the prospect of potentially earning £10,000.

Table 3 Propensities to bank

	20	50	100	200	300
Round 1	0.051	0.186	0.426	0.545	0.587
Round 2	0.073	0.295	0.510	0.736	
Round 3	0.108	0.321	0.611	0.732	
Round 4	0.143	0.435	0.670		
Round 5	0.161	0.507	0.800		
Round 6	0.251	0.537	0.756		
Round 7	0.314	0.605			
Round 8	0.650	0.732			

if the probability of a correct answer is lower, the probability that the money is lost will be higher, so one could expect the number of banks to increase. The probability of a correct answer depends on two factors: the difficulty of the question, and the general knowledge of the contestant. Therefore, one could expect the number of banks to increase when questions are more difficult, and when a contestant has less general knowledge.

In Table 2, we saw that, over rounds, questions become more difficult to answer. Hence, we expect that the propensity to bank at a given amount on the chain, is higher in later rounds. Table 3 gives a breakdown of the propensities to bank, conditional on both the round and the amount of money on the chain. For example, the top-left entry indicates that there was a bank in 5.1% of all the cases where the chain was at £20 in round 1. In the table, we have only included cells with at least 30 observations.

Conditional on the round, contestants do indeed bank more often when the amount on the chain is higher. In every single row, the propensity to bank is increasing with the amount of money at stake. Also, conditional on the amount of money on the chain, the probability of banking is almost always higher for later rounds. This establishes (a) and (b).

Given the limited number of questions many contestants receive, it is impossible to assess the state of knowledge of all contestants in our sample. However, (c) does imply that people that do not bank, are on average more knowledgeable. If this is true, we should have that *ceteris paribus*, there are less correct answers after a decision to bank, than there are after a decision not to bank.

Table 4 shows the relationship at the round level. Every column first gives the number of cases of a bank in that round, and then the number of cases of a no-bank, where a bank would have been possible. Next, the probabilities of a correct answer conditional on a bank and a no-bank are given. The last row tests for the statistical significance of the difference between the two probabilities, using a χ^2 test. It reports the p value on the null hypothesis of no difference.

In every single round, the probability of giving a correct answer is higher after a no-bank, which supports (c). Yet, using a χ^2 test, we find that the difference is significant at the 10% level only in rounds 2, 3, 6, and 7. In a test for the joint significance of the eight individual test statistics, we find a p value of 0.060. We have also tested our hypothesis on a more disaggregated level: conditional on both the round and

Table 4 Relation between banking and answering questions correctly

Round	1	2	3	4	5	6	7	8
Cases of bank	295	275	234	234	196	189	180	214
Cases of no bank	1158	941	739	609	495	389	296	128
Prob. correct answer when bank	0.766	0.658	0.585	0.598	0.536	0.460	0.433	0.467
Prob. correct answer when no bank	0.800	0.716	0.650	0.629	0.549	0.537	0.514	0.484
p value	0.204	0.064	0.076	0.412	0.743	0.082	0.090	0.759

the amount of money on the chain. This yields 20 cells for which we have sufficient observations to do a χ^2 test.[8] A χ^2 test for the joint significance of the 20 individual test statistics then yields a p value of 0.027. Thus, when looking at individual cells, we find a relation that is significant at the 5% level. This establishes (c).

5 A benchmark strategy

In the previous section, we established some comparative statics with respect to contestants' banking behavior. In this section, we go one step further, and test whether contestants use the banking strategy that maximizes expected total prize money. To do so, we choose an indirect approach.

Suppose that the contestants in a show always use an extremely naive strategy, by always banking in round i if and only if some amount x_i is on the chain. Clearly, this cannot be the optimal strategy. Players then do not use any information on their own knowledge, successes in earlier rounds, etc. However, if we are able to show that following such a strategy yields higher prize money than what contestants actually achieve on the show, we have unambiguously established that the strategy they use on aggregate, does not maximize expected winnings. But this is indeed the case. We find the following[9]

Result 2. *Total prize money could have been substantially higher on average by using the following benchmark strategy*:

(a) *bank at £200 during the first four rounds*,
(b) *bank at £20 during the last four rounds*,
(c) *bank at every opportunity during the last 15 s of any round*.

---

[8] Details of these tests are available from the authors upon request.

[9] This benchmark strategy was found by trying all possible strategies of this type for all the shows in our sample. That is, for round 1, we have calculated what total prize money over all shows would have been if contestants would have always banked if and only if the chain was at £20, what total prize money would have been if they would have always banked if and only if the chain was at £50, etc. In round 1, it turned out that over all shows the highest prize money would have been achieved if contestants would have always banked at £200. For the other rounds, we did the same exercise. Details are available from the authors upon request. The 15-s cut-off rule is somewhat arbitrary: the idea is that contestants should always bank when time is about to run out, since if they do not have enough time to answer the question, the money is lost. A 10-s cut-off does not yield qualitatively different results.

Table 5 Contestants' performance versus benchmark strategy

Round	1	2	3	4	5	6	7	8	Total
Money earned by contestants	638	418	311	246	159	141	123	101	2,339
Money earned with benchmark	648	455	317	267	203	175	144	116	2,555
Relative improvement (%)	2	9	2	9	28	24	17	15	9
Shows strictly beaten	39	42	38	39	64	53	53	38	58
Shows with equal performance	8	5	8	8	4	4	7	23	0
Shows strictly worse	30	30	31	30	9	10	17	16	19

Table 5 gives a summary of the relative performance of the benchmark strategy. We report on the money earned by contestants, and the money they could have earned using the benchmark strategy. The fourth row gives the number of shows in which the benchmark strategy does strictly better, the fifth row the number in which it does equally well, and the final row the number of shows in which it does strictly worse. Overall, using the benchmark strategy would have lead to a higher amount of prize money in 59 out of 77 shows. On average, our strategy would have earned the contestants £2,555, an improvement of 9% over their actual performance.[10]

One could object to our methodology by arguing that we use the benefit of hindsight in determining the best naive strategy ex post. Yet, we do not believe that such a criticism is valid. First, players on the show are able to use more information than we can in deriving our benchmark strategy, such as cues with respect to the general knowledge of their competitors, information of their performance in previous rounds, etc. Second, suppose we split our sample in two subsamples of roughly equal size. If we then use the same methodology to derive a benchmark strategy solely based on the observations in the first subsample, we arrive at the same benchmark strategy as mentioned above. Using that benchmark strategy, we can then improve upon contestants' earnings in the second subsample. In that sense, our methodology thus also improves earnings out-of-sample. Third, note that contestants have very little time to determine their banking strategy during a round. Hence, they could save time by deciding in advance to follow a very simple strategy, like the one we discuss here. This strengthens our results.

6 A simulation of the optimal banking strategy

6.1 Methodology

In the previous section, we established that contestants do not use the optimal (i.e., prize-money maximizing) banking strategy. In this section we present an algorithm that allows us to derive the optimal strategy, under some simplifying assumptions. We use that algorithm to obtain further insight into the difference between actual

---

[10] Note that the total amount does not equal the sum of the individual amounts of the previous rounds, since prize money in the last round is trebled.

and potential winnings, and, more importantly, to analyze how contestants' banking behavior deviates from optimal banking behavior.

Our algorithm proceeds as follows. Suppose that in some round, contestants will receive T questions. For simplicity, suppose that this value is given and known in advance. Assume that there still is time to bank after the last question has been answered. Questions are numbered $t = 1,...,T$. At any point, the state of the chain is denoted by s and the amount of money secured so far in this round is denoted by k. Thus, at any point, s is the number of uninterrupted correct answers that has been given before a contestant receives his question, and k is the amount of money that has already been secured in this round. We thus have $s \in \{1, \ldots, T\} \equiv S$, and $k \in \{20, 40, 50, 60, 70, \ldots, 1000\} \equiv K$. The amount that is added to the total when the contestant banks, is denoted $B(s)$. Thus $B(1) = 20$, $B(2) = 50$,... and $B(s) = 1000$ for $s \geq 9$. The ability of the player that receives question t is denoted p_t. This is the a priori probability that this contestant will know the correct answer to a question in that round.

The optimal strategy is solved using backward induction. First, consider question T. We can calculate the expected amount of money that will ultimately be won in this round if this contestant banks. That amount depends on the amount of money k secured in this round, the current state of the chain s, and the ability of the current player p_T. We denote this expected amount as $V_B(T, k, s, p_T)$. Similarly, the expected prize money that will ultimately be won in this round if the contestant does not bank, is denoted $V_N(T, k, s, p_T)$. We have

$$V_B(T, k, s, p_T) = (1 - p_T)\min\{k + B(s), 1000\}$$
$$+ p_T \min\{k + B(s) + B(1), 1000\}, \qquad (1)$$
$$V_N(T, k, s, p_T) = (1 - p_T)k + p_T \min\{k + B(s + 1), 1000\}. \qquad (2)$$

The first expression can be understood as follows. If this player banks, the total amount of money secured in this round moves from k to $k + B(s)$, provided that this amount does not exceed 1000. If she gives a wrong answer, which happens with probability $1 - p_T$, then this also equals total prize money in this round. If she gives a correct answer, which happens with probability p_T, there will be a final opportunity to bank after her answer, and the round total is $k + B(s) + B(1)$, again provided that this does not exceed 1000.

Now suppose this contestant does not bank. If she gives a correct answer, which happens with probability p_T, total prize money in this round will then be $k + B(s + 1)$, provided this does not exceed 1000. Otherwise, the total for this round will simply be k, hence (2). Conditional on the contestant making the optimal decision, the expected payoff can be denoted

$$V(T, k, s, p_T) = \max\{V_B(T, k, s, p_T), V_N(T, k, s, p_T)\}. \qquad (3)$$

Consider question $T - 1$. Suppose the state of the chain is s. If the player receiving a question banks, the expected amount of money that will be made during the round (including this bank), is given by

$$V_B(T-1,k,s,\boldsymbol{p}) = p_{T-1}V(T,\min\{k+B(s),1000\}1,p_T)$$
$$+ (1-p_{T-1})V(T,\min\{k+B(s),1000\},0,p_T),$$

where we write V_B as a function of all abilities, which we denote as $\boldsymbol{p}=\{p_1,\ldots,p_T\}$. This can be understood as follows. By banking, the amount of money secured increases to $k+B(s)$, provided this does not exceed 1000. The final contestant thus faces an amount of money secured in this round equal to $\min\{k+B(s),1000\}$. With probability p_{T-1}, contestant $T-1$ gives a correct answer, so the final contestant faces a state of the chain that equals 1. Otherwise, the final contestant faces a state of 0.

Now suppose contestant $T-1$ does not bank. The final contestant then faces an amount of money secured in this round that equals k. With probability p_{T-1}, contestant $T-1$ gives the correct answer, so the final contestant faces a state of the chain of $s+1$. Otherwise, contestant $T-1$ gives the wrong answer, so the final contestant faces state 0. Hence,

$$V_N(T-1,k,s,\boldsymbol{p}) = p_{T-1}V(T,k,s+1,p_T) + (1-p_{T-1})V(T,k,0,p_T).$$

If contestant T makes the optimal banking decision, we have

$$V(T-1,k,s,\boldsymbol{p}) = \max\{V_B(T-1,k,s,\boldsymbol{p}), V_N(T-1,k,s,\boldsymbol{p})\}.$$

Similar relations hold for any question t with $1 \le t < T$. We thus have the following recursive system:

$$V_B(t,k,s,\boldsymbol{p}) = p_t V(t+1,\min\{k+B(s),1000\}1,\boldsymbol{p})$$
$$+ (1-p_t)V(t+1,\min\{k+B(s),1000\},0,\boldsymbol{p}), \qquad (4)$$
$$V_N(t,k,s,\boldsymbol{p}) = p_t V(t+1,k,s+1,\boldsymbol{p}) + (1-p_t)V(t+1,k,0,\boldsymbol{p}), \qquad (5)$$
$$V(t,k,s,\boldsymbol{p}) = \max\{V_B(t,k,s,\boldsymbol{p}), V_N(t,k,s,\boldsymbol{p})\}. \qquad (6)$$

This system can be solved for any vector of known abilities $\boldsymbol{p}$.[11] Given such an ability vector, Eqs. (1)–(3) can be used to calculate $V(T,k,s,\boldsymbol{p})$ for all possible states (k,s), and given the ability vector $\boldsymbol{p}$. Given that matrix, the matrices $V(t,k,s,\boldsymbol{p})$ can then be calculated using the system (4)–(6), for all $t < T$ and again for all possible states (k,s). Hence, at any point in time, given the amount of money on the chain, the amount of money already secured, and the ability of all the contestants, this algorithm uses backward induction to determine whether or not it is an optimal strategy for the current contestant to bank. Of course, this does assume that all future contestants will bank optimally as well.

Our algorithm now proceeds as follows. First, we determine whether the answer of the first contestant is correct. We then move to the second player. If the first answer was correct, the second contestant uses the algorithm above to determine whether

---

[11] Note that, if there are two players left, one of them will receive all the even-numbered questions, and the other one all the odd-numbered questions, so we have $p_1 = p_3 = p_5 = \ldots$, and $p_2 = p_4 = p_6 = \ldots$. With three players we have $p_1 = p_4 = p_7 = \ldots$ etc.

or not to bank, by comparing $V_B(2, 0, 1, \boldsymbol{p})$ and $V_N(2, 0, 1, \boldsymbol{p})$. We then determine the answer of this second contestant. Suppose she did not bank and gives a correct answer. The third contestant then compares $V_B(3, 0, 2, \boldsymbol{p})$ and $V_N(3, 0, 2, \boldsymbol{p})$ in deciding whether to bank. If the second contestant did bank and gave the correct answer, the relevant comparison is that between $V_B(3, 20, 1, \boldsymbol{p})$ and $V_N(3, 20, 1, \boldsymbol{p})$. If the second contestant did not bank and did not give the correct answer, the comparison is that between $V_B(3, 0, 0, \boldsymbol{p})$ and $V_N(3, 0, 0, \boldsymbol{p})$. We then determine the answer of this third contestant. This process continues until the last question is posed, or the contestants have managed to bank £1000. For our simulation, we simply use the observed string of correct and wrong answers from our dataset, in the way that it occurded in the round and the show that we are simulating.

In the next subsection, we compare the expected earnings derived from our algorithm with the earnings of the contestants in our dataset. That analysis will also serve as a starting point for our analysis in Sect. 7, where we investigate the extent to which the contestants' behavior differs from the optimal banking strategy.

6.2 Comparing optimal strategy and contestants' performance

In this subsection, we compare the contestants' behavior with that prescribed by our algorithm. To do so, we need an estimate of the individual abilities of the contestants, which we use as an input for the algorithm. Denote the probability that player i gives a correct answer in round j of show k as p_{ijk}, with $i \in \{1, \ldots, 9\}$, $j \in \{1, \ldots, 8\}$, and $k \in \{1, \ldots, 77\}$. In our algorithm, we could use different abilities for each contestant in each round of each show, letting the p_{ijk}s differ over i, j, and k. Yet, it is impossible to obtain a reliable estimate for all those p_{ijk}s. Therefore, we assume that the abilities of all contestants within a round in a given show are equal, but that these abilities differ between shows, thus $p_{ijk} = p_{\cdot jk}$, for all i, j, k. Also, we assume that these abilities are equal to the observed fraction of correct answers in that show in that round.

We thus perform the following exercise. First, we look at round 1 in show 1. We calculate the fraction of correct answers in that round. Then, we run our algorithm, setting abilities of all players equal to that fraction, setting the number of questions equal to the observed number of questions, and using the observed sequence of correct and incorrect answers.[12] We repeat the analysis with round 2 of show 1, and with all other rounds in all other shows as well. This exercise yields for each round in each show the true winnings, and the winnings according to our algorithm. Finally, we compare the averages of these amounts over all shows. Table 6 summarizes our findings. We report the performance of the contestants in the show (the column labeled 'Cont'), and the relative improvement of the simulated optimal strategy against the actual contestants' performance.

---

[12] Note that there is a problem if contestants have reached the £1000 before time has run out. For those cases, we have estimated the total number of questions in that round as the number that would have been reached had contestants continued to answer questions in the same speed as they had done so far. Also, to allow for the cases in which contestants already reach £1000 after just nine questions, we put an upper bound of 0.95 on the ability estimate that we use. All these shortcuts, however, have no discernible effects on the results of our simulations.

Table 6 Comparing strategies

	Cont.	Optimal	Improve (%)
Round 1	638	+84	13
Round 2	418	+70	17
Round 3	311	+31	10
Round 4	246	+34	14
Round 5	159	+66	42
Round 6	141	+46	33
Round 7	123	+21	17
Round 8	101	+15	15
Total	2339	397	17

The results suggest that, by using the optimal banking strategy, contestants could on average have improved upon total winnings by 17%. Admittedly, by deriving this result, we do assume that players are able to correctly estimate the average ability of all contestants. We also made the simplifying assumption that, within a given round in a given show, all contestants have the same ability. In the real world, of course, abilities do differ. Therefore, we have also investigated the effect of increased heterogeneity in abilities.

To do so, we kept the average abilities in a given round in a given show constant, but increased the variance of those abilities.[13] In these simulations, we do not use the *actual* answers, as we did above, but rather *simulated answers*: for each question posed to contestant t, we draw a random number from a uniform distribution between 0 and 1. If this random number is less than or equal to p_t, then this represents a correct answer. In our simulations, we find that average earnings increase as the variance in contestants' abilities increases. This suggests that the results we report in Table 6 are a lower bound on the true increase in average earnings when using the optimal strategy. We thus have

Result 3. *By using the optimal banking strategy, contestants could have increased their total prize money by at least 17% on average.*

6.3 The propensity to bank

To gain further understanding of the optimal banking strategy, we look at the propensities to bank in different scenarios. To do so, we looked at a large number of simulations, again with simulated answers. We report the results in Table 7.

The column headings report the average ability p_{AV} of players in a simulation. Actual abilities are taken to be evenly distributed on $[p_{AV} - d, p_{AV} + d]$, where d is given by the number in the first column. The numbers in the cells then report the average propensities to bank for different states of the chain, as observed in the output

---

[13] Details are available upon request.

Table 7 Simulated propensities to bank

		0.45	0.50	0.55	0.60	0.65	0.70
0.00	20	1.000	1.000	1.000	1.000	0.042	0.000
	50	na	na	na	na	0.000	0.000
	100	na	na	na	na	0.001	0.001
	200	na	na	na	na	1.000	1.000
0.05	20	1.000	1.000	1.000	0.704	0.139	0.016
	50	na	na	na	0.000	0.010	0.001
	100	na	na	na	0.011	0.000	0.002
	200	na	na	na	1.000	1.000	0.708
0.10	20	1.000	1.000	0.954	0.598	0.264	0.077
	50	na	na	0.000	0.174	0.098	0.017
	100	na	na	0.000	0.132	0.092	0.009
	200	na	na	1.000	1.000	0.867	0.552
0.15	20	1.000	0.989	0.776	0.538	0.293	0.119
	50	na	0.000	0.449	0.275	0.146	0.050
	100	na	0.000	0.329	0.252	0.173	0.076
	200	na	1.000	1.000	0.971	0.749	0.496
0.20	20	1.000	0.826	0.630	0.510	0.296	0.147
	50	na	0.644	0.524	0.321	0.164	0.083
	100	na	0.474	0.411	0.306	0.199	0.105
	200	na	1.000	1.000	0.897	0.647	0.484

of the simulations. Consider, for example, the entry (0.60, 0.10). For this case, we performed 10,000 simulations with the following parameters: the number of contestants equals 7, the number of questions equals 24, and the abilities of players are evenly distributed on the interval [0.50, 0.70]. In these simulations, we observe a propensity to bank at 20 that equals 0.598. Hence, out of all the possibilities of banking at 20 in our simulations, we observe an actual bank in 59.8% of the cases.

Some patterns are worth mentioning. First, when abilities increase (or, equivalently, as questions become easier to answer), we see that in most rows and for most values of the amount of money on the chain, the propensity to bank decreases. This is also what we observe in the behavior of contestants. Second, the observed propensities to bank at 50 are virtually always *lower* than those at 20, different from part (c) of result 1. This is due to two effects.

First, consider a case in which abilities of contestants are relatively high. In that case, simulated contestants initially always bank at 200. However, once an amount of 800 has been reached, or when the end of the round draws near, they switch to a strategy of always banking at 20. Intuitively, this can be understood as follows. There is an option value of not banking: not only does the money chain reach a higher amount with a correct answer, but there also is the possibility of reaching even higher amounts. This option value is lower if the amount of money secured is relatively high, and if the number of remaining questions is low. In those cases, contestants have more of an

incentive to bank already at lower amounts. This implies that, for the aggregate data, we find a high propensity to bank at 20 and at 200, and a much lower propensity to bank at intermediate amounts.

The second reason why banking at 50 is observed less often in the simulations than banking at 20, is related to the way the money chain is designed. It is possible to show that, in a case where all contestants have the same ability, rounds go on indefinitely, and contestants use the banking strategy that maximizes their expected return per question, banking at 50 or banking at 100 is *never* an optimal strategy, regardless of the common ability that we assume.[14] This effect is still discernible in the simulation data, where the end-of-round and high-chain effects are taken into account.

Arguably, we therefore have that the propensity to bank at 50 and at 100 should be lower than the propensity to bank at 20. Note, however, that this is *not* what we observe in our real-world data. There, the observed propensities to bank are monotonely increasing in the amount of money on the chain, as we observed in Sect. 4. Apparently, contestants use an intuitively appealing rule that nevertheless is arguably not a part of an optimal strategy. We have thus established the following:

Result 4. *Consistently using a higher propensity to bank when the amount of money on the chain is higher, is* not *a part of the strategy that maximizes expected total prize money*.

7 Do contestants bank too early or too late?

We have established that contestants in *The Weakest Link* do not use the banking strategy that maximizes total prize money. If they would do so, they could increase their earnings by at least 17%. We have thus established that contestants bank in the 'wrong' manner. Yet, we do not know exactly in which way they are 'wrong'. In other words, do contestants bank 'too early' or do they bank 'too late'? In this section, we try to answer that question. Banking too late would suggest that contestants are either risk-loving or overconfident, in that they overestimate their ability of giving a correct answer. Banking too early would suggest that contestants are either risk-averse or underconfident.

We perform the following analysis. For every single banking decision in our dataset, we use our algorithm to derive the optimal banking decision, given the amount of money on the chain, the amount of money secured, the number of questions that remain, and our assumption that ability of all contestants in a given round in a given show are equal. Then, we compare our optimal banking decision with the decision taken by the contestant. This allows us to classify every single banking decision into one of four categories: correct banks (the contestant banks, and this is also the optimal strategy), incorrect banks (the contestant banks, but the optimal strategy is not to bank), correct non-banks, and incorrect non-banks. The results are given in Table 8. The column labeled CB gives the overall number of correct banks in round 1 in our 77

Table 8 Classification of banking decisions

	CB	CN	IB	IN	Over	Under	u/o
Round 1	77	1057	260	137	0.197	0.640	3.24
Round 2	60	838	267	136	0.242	0.694	2.87
Round 3	115	456	168	320	0.269	0.736	2.73
Round 4	131	325	150	320	0.316	0.710	2.25
Round 5	134	184	97	338	0.345	0.716	2.07
Round 6	148	128	85	281	0.399	0.655	1.64
Round 7	174	59	39	256	0.398	0.595	1.50
Round 8	191	39	83	98	0.680	0.339	0.50

shows, the column IB gives the number of incorrect banks, and CN and IN the number of correct and incorrect non-banks, respectively.

Consider a case where, according to our simulation, a contestant should have banked. If, in such a case, the contestant did not bank, then we can interpret this as a bank that is 'too late': the contestant should have banked already, but did not do so yet. Similarly, if there should be no bank, but the contestant did bank, then we have a bank that is 'too early': the contestant should not have banked yet, but did do so already. We define the fraction of overbanks as #IB/(#IB + #CN). It equals the fraction of cases where contestants make the wrong decision, and the right decision is not to bank. The fraction of underbanks is the fraction of cases where contestants make the wrong decision, and the right decision is to bank: #IN/(#IN + #CB). Column 6 in Table 8 gives the fraction of overbanks in each round, column 7 the fraction of underbanks, and column 8 the ratio of underbanks to overbanks.

Note that in rounds 1 through 7, the fraction of underbanks is overwhelmingly higher than the fraction of overbanks. In round 1, the difference is by a ratio of more than 3:1. This ratio monotonically decreases until round 7, where it is still 3:2. Yet, in round 8, the picture dramatically changes. The fraction of overbanks is now double the fraction of underbanks. We thus have the following:

Result 5. *During the first seven rounds, 'wrong' banking decisions are primarily decisions to bank too late. In round 8, 'wrong' banking decisions are primarily decisions to bank too early.*

Hence, in the first seven rounds, contestants often do not bank yet, when the optimal strategy prescribes that they should. In round 8, this changes entirely. Contestants now become much too conservative in banking. It is hard to believe that this is due to a sudden shift from risk-loving to risk-aversion. Yet, it is equally hard to believe that there is a sudden shift from overconfidence to underconfidence. Note however that round 8 is a round in which the two remaining contestants merely play in order to increase total winnings. After round 8, there is no voting stage. This suggests that contestants deliberately bank too late in earlier rounds, in an apparent attempt to convince their competitors that they are confident in their ability, and fearing that banking earlier would signal lower ability and hence a higher probability of being voted off. In round 8,

when signaling is no longer necessary, contestants bank far too early on average. This is inconsistent with overconfidence. It thus suggests that, indeed, signaling is the only explanation for underbanking in the earlier rounds.

A signaling explanation would also suggest that as contestants have better information about each other's abilities, it becomes less of a necessity to try to convince the other players of one's own ability. Hence, one would then expect the relative number of underbanks to decrease as the show progresses. This is exactly what we find.

An increasing body of literature in psychology and economics suggests that people are overconfident when evaluating their own abilities (see e.g., Camerer and Lovallo 1999 or Nöth and Weber 2003). Bénabou and Tirole (2002) argue that rational agents have an incentive to build self-confidence. One of the three reasons they list is that self-confidence has a signaling value: "[…] believing oneself to be of high ability […] makes it easier to convince others that one does have such qualities" (p. 877). Interestingly, contestants on *The Weakest Link* also seem to try to convince others that they are of high ability. Yet, our results suggest that they do not really believe themselves to be of such high ability, since they do switch to overbanking in round 8. Camerer and Lovallo (1999) predict that "in hierarchical tournaments where "winners" at one level advance to the next level […] overconfidence will get stronger and stronger as people advance." (p. 316) We find that the opposite is true: as contestants advance, the number of underbanks relative to overbanks *decreases* which, if anything, suggests a decrease in overconfidence.

Note that the attempt to try to convince other players that one's ability is higher than it really is, does not necessarily have an effect on the other players' perception of this contestant's ability. Indeed, rational players will take these incentives to signal into account and discount them accordingly (see e.g., Mailath 1987). For the contestants, such signaling is then a prisoners' dilemma: they would all be better off if they could commit not to signal in this manner.

8 Conclusion

In this article, we studied the BBC game show *The Weakest Link* as a field experiment in decision making. We used the banking decision in that game, where a contestant chooses whether to secure a certain amount of money for her team, or to risk that money on a general knowledge question. In the latter case, should the contestant give the correct answer, the amount of money available increases substantially. We tested to what extent contestants make the optimal banking decision.

We established the following results. Contestants on this show are more likely to bank when the amount of money at stake is higher, when questions are more difficult to answer, and when a player has more general knowledge. These are simple rules that seem to make intuitive sense. Yet, we also showed that a strategy that maximizes the winner's expected earnings will *not* yield such a simple monotone relation between the amount of money at stake and the propensity to bank. Thus, although contestants seem to obey such rules when making their banking decision, some of these rules do violate the optimal strategy.

We also showed that, as a group, contestants do not use the optimal banking strategy, i.e. the banking strategy that maximizes expected total prize money. Using a computer algorithm to derive the true optimal banking, we found that players' average winnings could have been at least 17% higher or £397 for the ultimate winner. This is just a lower bound. We assumed players in a given round in a given show to be of equal ability. An increase in the variance of those abilities further increases expected average prize money.

In the first seven rounds, contestants systematically bank too late. The fraction of underbanks is far higher than the fraction of overbanks. Yet, the ratio between the two decreases over time. In round 8, after which no voting decisions are made, this picture changes entirely. The fraction of overbanks now is far higher than the fraction of underbanks. These observations seem consistent with the following story. Contestants deliberately bank too late, in an attempt to try to convince the other players that they are confident in their ability to do well on the show. This incentive becomes less prominent as the show progresses and contestants learn more about each other's abilities. In round 8, after which there is no voting decision, risk aversion becomes the dominant consideration, leading contestants to over-rather than to underbank. Our results suggest that contestants do not suffer from overconfidence. Yet, they do try to convince others that they are better than they really are.

The mechanism we identified in this article may also be present in other situations with uncertainty regarding agents' capabilities. Consider for example the case of a young economist who considers a suitable outlet for his work. First suppose that his department (and/or potential future employers) will only be able to observe his actual publication record. The choice of journal will then be a trade-off between the probability of acceptance and the quality of the journal. The interests of the economist and his department are perfectly aligned. But now consider the case in which the department can also observe where this economist submits his paper. In that case, he will be inclined to send his paper to a better journal than he otherwise would, in an attempt to signal that his quality is higher than it really is. The frequent rejections that will result are suboptimal for the department as a whole. Again, just as we observed in *The Weakest Link*, even though the incentives of the individual and the group seem perfectly aligned, the desire to signal leads to a decision that is riskier than the first-best solution.

As a final example, consider a firm that consists of a number of departments. The manager of each department proposes and implements a risky project. The probability that a project is successful, depends on the unobservable quality of a manager. The remuneration and career opportunities of a manager implicitly depend on his perceived ability. When the executive board can only observe whether or not a project is successful, the interests of the firm as a whole and each individual manager are perfectly aligned, and the manager will choose the project that maximizes expected profits. But when the executive board can also observe the projects that are implemented, a manager will be inclined to implement a riskier project than he otherwise would, in an attempt to signal that his quality is higher than it really is. Also in this case, even though the incentives of the individual and the group are perfectly aligned, the desire to signal leads to a decision that is more risky than the first-best solution. We have formalized this idea in Haan et al. (2008).

On the basis of our analysis, it is hard to judge whether contestants on *The Weakest Link* behave rationally. On the one hand, they fail spectacularly to use the banking strategy that maximizes total expected prize money. But on the other hand, our analysis suggests that this may at least partly be due to the fact that they deliberately bank too late, which may well be an equilibrium strategy in a signaling game where contestants sacrifice prize money to try to convince other players that they have a high ability to answer general knowledge questions.

References

Antonovics, K., Arcidiacono, P., & Walsh, R. (2005). Games and discrimination: Lessons from the weakest link. *Journal of Human Resources, 40*(4), 918–947.

Beetsma, R. M. W. J., & Schotman, P. C. (2001). Measuring risk attitudes in a natural experiment: Data from the television game show Lingo. *Economic Journal, 111*(474), 821–848.

Bénabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics, 117*(3), 871–916.

Berk, J. B., Hughson, E., & Vandezande, K. (1996). The price is right, but are the bids? An investigation of rational decision theory. *American Economic Review, 86*(4), 954–970.

Camerer, C., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review, 89*(1), 306–318.

Février, P., & Linnemer, L. (2006). Equilibrium selection: Payoff or risk dominance? The case of the weakest link. *Journal of Economic Behavior and Organization, 60*(2), 164–181.

Gertner, R. (1993). Game shows and economic behavior: Risk-taking on card sharks. *Quarterly Journal of Economics, 108*(2), 507–521.

Haan, M. A., Los, B., & Riyanto, Y. E. (2008). *Harmful monitoring*. The Netherlands: Mimeo, University of Groningen.

Levitt, S. (2004). Testing theories of discrimination: Evidence from weakest link. *Journal of Law and Economics, 47*(2), 431–452.

Mailath, G. J. (1987). Incentive compatibility in signaling games with a continuum of types. *Econometrica, 55*(6), 1349–1365.

Metrick, A. (1995). A natural experiment in Jeopardy!. *American Economic Review, 85*(1), 240–253.

Nöth, M., & Weber, M. (2003). Information aggregation with random ordering: Cascades and overconfidence. *Economic Journal, 113*(1), 166–189.

www.manaraa.com