Towards testing the theory of "On Committees of Experts", setting up an experiment

By Sander Renes

Towards testing the theory of “On Committees of Experts”,
setting up an experiment

Master of Science Dissertation

Abstract
In their 2007 paper “On Committees of Experts”, Bauke Visser and Otto Swank propose a theory about the decision making of committees of experts. So far this theory and the level of rationality it assumes have not been tested. In this thesis I propose a laboratory experiment to test this theory. In addition I propose several alternative experiments that could be used to test parts of the theory and results, as well as report some observations on the first actual trials of one of the alternative experiments.

Keywords:
Experiment, reputation, committees, experts, entrapment, escalation, decision making,

Thesis supervisor: Prof. dr. B. Visser
Erasmus University Rotterdam
Kamer H08-20
Postbus 1738
3000 DR Rotterdam
bvisser@ese.eur.nl

Author: Sander Renes
Willem Buytewechstraat 169B
3024 XG Rotterdam
sander.renes@gmail.com

Exam number: 295306

Table of contents

1. Introduction

A society without specialization in production seems unthinkable in these modern times, and without question it has brought us a standard of living and wealth previously unthinkable. This specialization has spread to all parts of our production structure. As we have specialized people for every product we produce, so we have experts for practically each type decision that needs to be taken.

In normal day-to-day life we trust these experts and we have good reason to do so, many of them have specialized just to be able to assist us in their field of expertise. Unfortunately, there are also many cases known where expertdom might actually cause problems. Examples include situations as diverse as doctors that cannot reach an agreement on what treatment is better for specific ailments, as well as the famous stock monkey that keeps beating the stock-analysts in picking the right stock. In both instances groups of experts, people who are supposed to know their field better than the rest of us, seem to do something wrong.

If I have a seizure I want the best possible treatment, it is hard for me to imagine that the best treatment for my seizure in England is not the best treatment for my seizure France. Just as it is quite stunning that a monkey picking bananas, seems to do a better job in prediction stock returns than many stock advisors. The question in both cases is, how can this happen? How is it possible that this group of experts cannot get it right?

In this thesis I will propose an experiment aimed at testing the predictions of an economic theory of expert group decision making, proposed by Visser & Swank (2007). I will start with a short summary of existing literature on decision errors similar to the one I am looking at, both at the individual and at the group level.

Even though a lot of research has been done by social psychologists, my focus will be on economic research. After this literature review I will focus more in depth on the actual theory of Visser & Swank (2007). This is followed by the introduction of the main experiment and several alternatives to this experiment. Before I round up with some concluding remarks, I will report on some trials I did with one of the alternative experiments and discuss what I learned from observing those trials.

2. Background

2.1 The existence of the decision error

Visser & Swank (2007) already refer to the survey of Gerling et al. (2005) to summarize large branches economic theory relevant to committee decision making. Since this summary is complete enough for me for the strands it summarizes, I prefer to focus on some other strands of literature that are akin to the problem I want to test in the experiment. Closely related to the problem is the literature discussing; entrapment, escalation error and sunk cost error and therefore I will focus on discussing the papers from these strands that I find important for my research.

In this ever more complex world, a lot of authors have defended the vision that group decision making leads to better outcomes. They believe this kind of decision making can rely on more information and different views, and thus should reflect more sides of the decision at hand.

Janis (1972) is one of the first to recognize that there can be negative side effects of group decisions. In his groupthink hypothesis he defends the view, that internal group pressures aimed at keeping the group together will cause the members to withhold information and be less critical than they would normally be.

Another possible source of errors is known under many names in different fields, although there are subtle differences between them, the overlap is more than substantial. So far I have been able to track it down as (examples between brackets); entrapment or escalation error (Brockner & Rubin 1985, Staw 1976, Staw and Fox 1977), sunk cost error (Kanodia, Bushman & Dickhaut 1989, Harrison & Harrell 1993 &1994), Concorde effect (Arkes & Ayton, 1999). This list is far from complete, but it gives an indication of the different schools that have published on the subject.
The effect these strands of literature study, focuses on the idea that people are less willing to give up projects they have invested in already because of the investments made before, not because of future profits. The link between these strands and the problem type seen in this paper is in the nature of the problem and the possible motivation for the “error” made. In short, most of the literature below looks at situation where people do not take the decision that maximizes the payoff to some project (abandoning it), but take an option with could have an

alternative or extra benefit (continuing). One of the stronger explanations for the persistence of the escalation error is face saving, otherwise known as self-representation concerns. A much mentioned, but in my opinion not necessary, aspect is the negative feedback criterion. In much of the older literature, entrapment is defined as continuing investment after negative feedback. In the newer papers the definition of escalation error is more refined, escalation is only considered an error if someone escalates beyond the point of profitability or narrow economic rationality. This last definition is more precise and complete than the requirement of negative feedback, since it includes any project that is not economically worthwhile, no matter the feedback.

The Staw (1976) business-case experiment has set the standard for much of the later literature. In this experiment the subjects are asked to select one of two divisions to receive the ten million in R&D money. Half of the subjects then gets information that sais the division they have chosen failed after the extra boost, the other half gets the information that the division they have chosen did well with the money they gave it.
Subjects are then asked to spread 20 million of research funds between the two divisions of the company. The main dependent variable is the amount they invest in the failing division. Staw finds that subjects that get negative feedback have a greater tendency to invest in the division they have originally chosen..

Teger (1980) differentiates between males and females, by having subjects participate in an entrapment experiment, an all-pay auction for one dollar. His surprising outcomes indicate that male subjects get more entrapped individually than females individually, but they decrease the amount of entrapment by half when they are teamed up in pairs. Same sex teams of females, however, show a greater tendency to get entrapped than female individuals.

Baron & Byrne (1981) notice a tendency for groups to take the more extreme decision, this tendency has been known as the risky shift since. This result has been nuanced in some of the entrapment literature.

Bazerman, Giuliano & Appelman (1984), for instance, use the setup of Staw (1976) and find that if you put people in entrapment situations, people that feel more responsible get more entrapped. They also find that groups get less entrapped than their peers that take the same decision individually. They explain this group effect by the possibility groups offer to dilute, or reduce the amount of responsibility one feels. It is important to notice that if one feels responsible, that means the decision somehow reflects back upon oneself.

A clear critique to this paper is given by Moon et al. (2003), when they discuss the differences between individual and group decisions. They argue that, by giving the problems to the subjects only after creating the committees, Bazerman et al. (1984) create a low responsibility setting, since the subjects do not have to defend their first reaction as an individual. To prove their point, they differentiate between individual and group commitment by manipulating the moment subjects get the information about the project. They use students in a management course as subjects.

They find that, committees that have time to evaluate as individuals, which is akin to the private signal that follows from Visser & Swank (2007), show significantly less abandoning of the project and more instrumentalism than committees that receive the initial information as a group. Decision certainty, however, is lower in the case where the individuals form an opinion prior to being in the committee than in pure committee decision making, which probably leads to incrementalism they find and through this to some consensus forming.

The individual decision makers also show significantly less abandoning than the pure committee decision makers, but more than the committees with private opinion forming. So the most rational responses are found in the pure committee decisions, followed by the individual decisions, and last come the mixed decisions.

In a second study they ask subjects to make a budget decision over different options. The committees that take a decision purely as a committee again show less investment in projects with a negative feedback (despite their 'responsibility') and more risk taking i.e. increased investment in projects that have been going well upon till then.

In the main experiment I propose subjects are not required to do anything with or for their signal until the committees are formed, this makes it unlikely that they perceive the signal as relating to them. Therefore they are unlikely to feel extra motivation to defend it.

2.2 The reasons for entrapment

In much of the older literature on entrapment (see for an overview Brockner & Rubin, 1985) the main explanation for this error is rationalization; a desire to get something from the earlier investments despite new information. But recently more focus has been applied to the possible (economically) rational component of this kind of decisions. Explanations entail, amongst others, project completion (Wilson and Zhang, 1997), risk/loss attitudes (Wong and

Kwong, 2007), a desire to “save face” (Brockner, Rubin & Lang, 1981) and other self-representation or reputational concerns (Visser & Swank 2007) or future concerns (Kanodia, Bushman & Dickhaut 1989). In most of the experimental literature this translates to a decision on a project that has a dual payoff structure, on the one hand there is the project value that can be maximized and on the other hand there is another value to the decision maker. One of the hardest things for all researchers, especially the economic researchers, is to interpret this second part of the payoff. In this overview I focuses on the overlap between the papers and mostly on the reputation, self-representation or image concerns explanation.

The search for explanations for this effect, however, has not stood still. The original rationalization approach has been under attack for some time, but has also been strongly defended by its proponents, see Brockner (1992). An important observation is made by Conlon and Garland (1993); they claimed that many of the effects found in the entrapment research can be explained by another effect, the percentage of completion effect. In most of the literature based on the setup of Staw (1976) the subjects in the experiments do not know how far along the projects are, they only know how much is invested in the projects. In their experiment Conlon and Garland introduce a new business case, again with a decision whether or not to further fund an R&D project. The negative feedback in this scenario is given in the form of information about the progress a close competitor is making (he is very close and better and cheaper vs. he is not there by long). Extra information is provided about the completion rates of the R&D project and information about how much has been spent already. By taking into account the percentage of completion of the R&D projects the subjects can invest in, Conlon and Garland make the sunk cost effect disappear totally if the two extreme groups were compared. There is a small problem with the intermediate groups, they show no particular inclination to either of the two extremes.

There are however two problems to this setup. First, it is quite different from the traditional Staw (1976) setup, especially in the negative feedback aspect, so it is hard to check how the subjects might have reacted to the manipulation without the completion information. Second, it is even harder to transpose this idea, let alone the result, to the non-economic examples of the entrapment effect, since project completion is a difficult concept to impose on some of the experiments.

On the other hand, the experiment does show that completion-concern can be (part of) the driving force behind many of the results found. Combined with risk aversion and ambiguity aversion, the continuing of projects could maybe be explained in a manner much more rational than thought before, since starting over does deliver a lot of uncertainty.

This relatively new tendency to search for more future related and more rational concerns and approaches is seen very strongly in a new strand of research. Tan & Yates (1995), for instance, review and combine research into the prohibitive aspects of formal instruction on how to handle sunk cost. They compared business students with non-business students and had them choose between options in two scenarios. In the ski-trip scenario the first option has a higher sunk cost, the other is expected to be more fun. In the investment scenario money either has been invested or not to begin with and then subjects are asked whether they want to invest more. They have mixed results, the business students do seem to pick up on the sunk cost in the investments scenario, but not in the ski-trip scenario. The non-business students miss the sunk costs in both scenario's. In a second study, besides dropping the ski-trip scenario, they add information on expected returns of the investment and the difference between the group with and without prior-investment disappeared for the non-business students as well, they no longer have a significant sunk cost effect.

Fear of regret has also been named as one of the more rational driving forces of entrapment. Wong and Kwong (2007) show that escalation increases if the subjects face the possibility of more regret by manipulating the amount of information people receive about the outcome of their decision. In their setup half of the subjects gets to know the outcome of escalation even if they do not escalate, this half escalates significantly more than the other half. They explain this effect by pointing at studies in the social psychology that show that we regret more that what we have not done, but should have done, than what we have done, but should not have done.

An earlier attempt to find rational explanation for the sunk cost error is given by Kanodia, Bushman & Dickhaut (1989). They use an agency theoretic approach to show that a manager/entrepreneur with career concerns could choose to continue a failing project. In their model a manager is either able or non-able and he picks a project to invest in, in period one. There are two projects and two possible states, for each of the two states only one project is optimal. An able agent has a better chance of finding out what state he is in. After a while the entrepreneur receives private information about the success of the project, if the wrong project has been chosen, switching is still profitable. They find that the managers have a tendency to continue the initial project to hide from the outside world the fact that he has picked the wrong project initially. This way he signals to the outside world that he is able, and is therefore offered a better contract in the second period.

2.3 Visser & Swank (2007)

Visser & Swank (2007), like Kanodia et al (1989), use agency theory in a binary state, binary ability and binary choice model. Since, they approach this problem from group decision theories like the herding and cascading literature, see Gerling et al (2005) for an overview, some things are different from the research mentioned above.

They start with a committee of decision makers, each of which is either smart or not, but the decision makers do not know which type they are. The prior probability that they are smart is common knowledge. Each committee member receives a signal about the project they have to decide upon. Smart members get a fully informative signal (success or failure), dumb members receive a signal selected at random.

The prior expected value of the decision, before the signals are received, is negative, so without further information the project should not be implemented.

The decision is taken in two stages, first the committee members discuss (send messages) and talk about the signals they have received and afterwards they vote. Both the sending of the messages and the voting is assumed to occur simultaneously, so there is no influence of the order of speech.

The members of the committee care about two things, first they care about the project payoff, the success it has, and second they care about their reputation (different experts can attach different weights to their reputation).

Reputation is modeled by the posterior probability the public attaches to the expert's ability (the chance he is smart). To make this assessment the public only has information on the prior probability, the ex ante expected value of the project and the decision the committee takes.

The model predicts that experts, defined as people that care for their reputation, will want to talk with one voice, i.e. all claim to have had the same idea and agree full heartedly with the choice made. Furthermore there is a tendency for the committees to take the more unlikely option, because this signals ability. This is an effect of the ex ante negative value of this option to the experts. To maximize investment value, they have to be relatively sure about the investment, so if the signals are mixed it is most likely the do not invest. This means that if they invest, there is a relatively large chance the signals concurred. Furthermore, since the chance the experts get the right signal is supposed to be above 50% (they are experts), the chance the signals concur because both received the right signal is larger than the chance they both received the wrong signal.

There are four difference between this setup and the one in Kanodia et al.(1989) that I think are important. First, it starts from committees of decision makers and not individuals. Second, the past is not relevant; whilst in most of the other literature the past was essential, this model only has one period.
Third, where Kanodia et al. (1989) create career concerns, equal to the gain in wage from appearing able, Visser & Swank (2007) proposes reputational concerns. This is a more general interpretation of the probability update, which makes it possible to use this update in more situations. Important to notice is also the choice the experts make, where it is just binary with Kanodia et al., it is important that there is one unlikely option, an option that has an ex ante lower expected value than the other in Visser & Swank.

The theory gives several assumptions and outcomes that are worth testing. First, it seems important to test the effect this dual payoff structure has on agents. It seems logical that they are willing to sacrifice some of one type of payoff for more of the other type, but how will they do that?
Second, there is the higher layer of rationality. It has been proven in much of the individual decision making literature (Kahneman & Tversky, 1972 a.o.) that people have problems with the Bayesian update. The experts in Visser & Swank (2007) do not only perform an update, they anticipate on the update of the public when choosing their action. This level of rationality has not been tested before. This also makes it interesting to see how the public reacts to this anticipating agent. Do they trust him because of, or maybe despite his actions? or do they focus on something else entirely?

3. The main experiment

To test the theory I want to stay as close as possible to the theory of Visser & Swank (2007). Throughout the description of the experiment below I make the link between the setup en the theory clear by showing what specific part of the theory I am trying to operationalize.

3.1 The experiment

Al subjects start with a folder containing instructions about the experiment. All subjects get the information stating that some subjects will be committee members working for the investment bureau of some company, others will be monitors that are trying to buy the services of good analysts and that they retain this role throughout the experiment.

Next is an explanation of how the committee members are linked to (fictitious) analysts (random) in the company/experiment and that the committee members will try to "sell" their analyst to the monitor at the end of each round and the procedure for this sale.

The fictitious company will take the role of principal, the entity that cares only about the project payoff. The committee members are the experts in the model.

The experts' payment will reflect the project outcomes/the principle's demands, as well as a possible reputation bonus. They all receive a show up fee, 100 ECU, with which they can invest (or not).

For every 2 or 4 committee members another subject should be appointed as monitor. This monitor has a dual task, he checks the setup of the randomizers/state contingent signals (if possible), and he is used to measure reputation formation, i.e. he takes the role of the public, which is the entity to which the experts want to show that they are able.

Al committee members get the analysis of "their" analyst as a proxy for their signal.

The analysis should provide them a signal about whether the project should be undertaken from the project-value point of view (the principals' point of view). This can be done by giving them a simple yes or no signal.

The analysts are drawn from a pool (work for the company), it is known internally (and by the subjects) that about 50% of the analysts are actually good. If they have a good analyst they can trust him/her blindly, the bad analysts just advices to invest 50% of the times, independent of the project in question, since he knows that is about the success rate. The analyst the subjects are randomly paired up with is "new to them", so they do not know whether or not it is a good one.

The committee members are then paired up at random and given the opportunity to invest as a committee. All committees will be presented with an investment opportunity drawn from a larger set. Of the total set of investment projects 50% will fail, this should be known by the subjects. This 50% success/failure rate means that the payoff of investing is -20 ECU with a chance 0.5 and + 10 ECU with chance 0.5.

This creates an ex ante negative expected value (.5*-20+0.5*10=-5), like p is in the model. These values can change as long as the expected value, before additional information is received, is negative.

After they are paired up in committees, the committee members get to 'talk', i.e. send a message to their fellow committee member about the signal they received. The theory predicts that in this setup, where both members care equally about their reputation, honesty is the equilibrium strategy. There are two possible situations; they received identical signals (y,y; n,n), or they did not receive identical signals (n,y; y,n), they should update accordingly. The updates are performed in appendix A.

After the deliberation stage comes the voting stage. The decision to invest should be taken unanimously, if one of the two members votes against the investment, it will not be made. Both the deliberation and the voting is private, the monitor will only know the outcome of the vote.

Both committee members then individually send a message to the monitor. This message is limited to a statement about the state they think they are in; they can only say: "I think investing is a good idea" or "I think investing is a bad idea". The committee members are free to select the message they want, as long as they know the monitor can use the information. Note that, due to the payoff structure, all committee members want to sell their analyst, so any communication about the sale will be uninformative anyway.

Following Visser & Swank (2007) the reputational bonus should be earned by means that can be, but do not have to be conflicting with the desire to maximize investment value. This is achieved by giving the monitor the opportunity to buy the analysts of the committee members. After the monitor sees the outcome of the investment decision and receives the messages the committee members send him, he has to decide on whether or not he wants to buy any analysts, or maybe both of them. If the committee members can convince the monitor their expert is good, he should be willing to buy. So keeping a good reputation with the monitor increases the payoff of the committee members.

To the monitor buying an specific analyst costs 10 ECU, so buying both analysts costs 20 ECU. He will receive 19 ECU for each smart analyst he buys. The values are chosen in such a way that the monitor only needs to have a small indication that the analysts are better than average (ex ante) for him to be inclined to buy an analyst. Furthermore, the ex ante expected value should be below the value of a fair bet to ensure that not all analysts are bought, since the posterior probabilities of having a good analyst should be near 0.5 (see appendix).

Since this experiment is about measuring the effect reputation forming has on the decision making, we let it enter in the payoff schemes of the subjects. Besides the project payoff, the committee members get the price paid for their analyst times a multiplicative factor (λ). Half of the sessions will constitute the control group with $\lambda=0$, for the other half, the reputation group, $\lambda=0.25$.

In summary this setup follows the theory as close as possible. Like in Visser & Swank (2007) all the choices that subjects have to make are binary, to invest or not to invest in a project, to buy or not to buy. The state and the ability of the experts are made binary as well. The pay-off structure of the model is followed as closely as possible. The committee members get the project payoff + the (reputation) price paid for their analyst times λ, that proxies for the weight attached to the reputation. The monitor pays 10 ECU per analyst he buys and gets 19 ECU per good analyst bought and a show up fee for being there and checking whether or not we tell the truth about the probabilities (he can check the randomizer used to produce the signals and the state etc.). It might be necessary to pay the monitors an extra fee since they are expected to average out during the experiment, while the committee members earn extra money by clever investment and selling their analysts.

By using a subject as monitor we can also get the subjects to believe that we are fair, because the one that hands out the bonus does not know the state (he checks the setup, not the outcome) and thus, that they can influence him toward believing them.

3.2 Equilibrium behavior

The split-up in two groups by λ should mean two types of behavior in the committees, see Visser & Swank (2007). For the control group reputation plays is irrelevant, they should invest if and only if they receive two positive signals. They have no reason to send a message that is different from the analysis they receive, I expect honesty.

The reputation group on the other hand will want to influence the monitor towards buying their analyst. Because of this ulterior motive, Visser & Swank (2007) predict they will invest in some projects (fraction β) that have conflicting signals and should always speak in one voice, that is, claim to have received the signal yes (no) if the investment is (not) made. This is because in equilibrium investing should be an indication of signal concurrence, and thus of ability.

Remember that the payoff of investing was -20 if it goes wrong and +10 if the investment succeeds and for the reputation group λ=0.25. This has two implications for this experiment, first the decision to invest in case of mixed signals has to have zero expected value, since that is also the expected value of not investing (mixed equilibrium). Second for the bank/monitor to buy an analyst, the expected value of doing so must be at least zero. If we use those restrictions, the equilibrium is solved with a beta of approximately 0.31. This beta leads to the set of posterior probabilities:

$\hat{\pi}(X = 1; \beta)$=0,526315779

$\hat{\pi}(X = 0; \beta)$=0,480000006

For the monitor this means that buying an analyst if X=1 has an expected payoff of 0,526315779*19-10=0,026239, while buying one analyst if X=0 would yields him 0,480000006*19-10=-0,89497. Since the expected value of the second analyst is identical to the first, the monitor should always either buy both or non. In equilibrium the monitor buys both analysts if and only if he sees that the investment is made (and receives the appropriate messages from the committee members). If the monitor buys only one analyst that is probably a sign of risk aversion, since buying two analysts basically doubles the value of the bet.

3.3 Alternative hypothesis

Visser & Swank (2007) suggest that experts will want to take the unlikely action to appear capable, so they take unnecessary risks. Intuitively one would expect something else. The decision the committee takes is a risky decision, if it is important some public or monitor trusts you, you want to appear cautious. In the experiment, that would translate to subjects in the reputation treatment that are at least as careful as the control group. It seems unlikely that agents that want to sell their analyst, will admit they disregarded his analysis, so the speaking in one voice will most likely remain under the alternative hypothesis.

If I assume the public likes careful experts and analysts, I expect that analysts are not only sold when the investment is made, but also if the investment is not made, just as long as the messages concur. Note that this means that the public does not use a Bayesian update to look at the probable ability of the analysts and the committee members anticipate this behavior. If significantly more analysts are sold after the investment is made, compared to without investment, that means the public updates in a manner close to Bayesian.

There is one slightly problematic aspect to this setup, if the monitors prove to be unable to see that investment is a sign of signal concurrence within the committee, they will not show a tendency to buy after investment. This means it might be hard to separate the alternative hypothesis from a missing update. However if significantly more analysts are bought if there is no investment, it is clear that the alternative hypothesis holds.

So in short what I expect to find can be summarized as:

Treatment	Control	Reputation		
Payoff	Outcome investment	Investment + 5*price analyst		
Expectations	First best behavior	Visser & Swank (2007)	Alternative hypothesis	Missing update
Investment	Only if y,y	y,y and some y,n	Only y,y	Only y,y
Messages	Truthful/ unimportant	Speaking with one voice	Speaking with one voice	Speaking with one voice
Analyst bought	More analysts are bought if the messages are identical and fit the investment decision	More analysts are bought if messages are identical and investment is made	More analysts are bought if messages are identical and no investment is made.	No significant difference between the amount of analysts bought whether the investment is made or not.

3.4 Extra checks

Manipulation check questions afterwards could check what strategy the monitor used to determine how many and which analysts to buy, as well as why the committees took the decision they took. This could help find out where things went astray, if the outcomes are not as predicted by theory. It can also help to evaluate whether or not the monitors spotted that investment by the committees was a sign of signal congruence and thus signaled a higher likelihood of having capable analysts.

Asking the committee members how they think the monitor judges the reliability of their analysts could also be used to see if they spotted the opportunity to influence the monitor and create a larger chance of selling their analysts. Basically, it is interesting to check for deliberate strategies.

4. Alternatives

The experiment mentioned above tries to test the entire model in one go, this seems a big task and I fear that it might prove too big of a task. That makes it interesting to test parts of the theory in separate experiments to find out where the reality is different from theory, as well as to change the setup of the experiment slightly to reflect some aspects of reality more clearly. I will split these alternatives up into two main categories. The first category, alternative approaches, comprises of experiments that are more or less the same as the main experiment. I will only explain the changes I want to make and why I think they might be beneficial.

The second category consists of alternative experiments, these are different experiments altogether, designed to check parts of the theory, as well as the level of rationality needed to be able to anticipate upon a (preferably Bayesian) update of the one responding to your actions.

4.1 Alternative approaches

Costly messaging

The model revolves around cheap talk, but the possibility to influence the monitor should have value to the committee members if their reputation matters. Making the message the committee members send to the monitor costly, would make it possible to show whether or not the committee members use a deliberate messaging strategy. It will show if the subjects pick up on the possibility to influence the monitor trough the message. If they do, I expect committees that are in the control group not to send messages, while the committee members that are in the reputation group should still send messages. This way the experiment would allow testing of one more hypothesis: if the control group sends significantly less (costly) messages than the reputation group the messaging strategy is deliberate.

Changing the signal structure

Instead of presenting the committee members with a simple yes/no analysis we could present subjects with profit prognoses, an internal rate of return and other financial data. This setup requires more explanation from the experimenters, but is more in line with existing research into the sunk cost effect, which makes the outcomes easier to compare. It also allows for more analysis because it makes it possible to change the gravity of distorting the decision (creates a p on an individual project basis) and thus see how far subjects are willing to go to get the other type of pay-off. The project payoff should then linearly depend on the difference between the expected rate of return and the required rate of return, so the subjects can easily compare outcomes. This should have the effect that some projects (small negative difference) are more easily accepted than some other projects (large negative difference), as long as the subjects believe investing boosts their reputation.

The difference in signal accuracy could be maintained by having some analysts just giving the wrong rates of return: “sometimes an analyst does not have a clue and just picks a value”. This makes the analysis more complex for the subjects and it requires more communication, so it will complicate the experiment. It will become more relevant to see whether the signals received are in line with the ones reported to the other project member.

An additional advantage to this setup could be that the subjects have to exert more effort to create their own individual opinion. This means they identify more with their signal, see Moon et al. (2003), in the same way that we expect experts to identify with their opinions. Although this is not directly the same ability update Visser & Swank (2007) uses as reputation, it quite clearly is a self-representation concern and therefore could easily be considered a reputational concern. This would essentially test the same hypotheses as the main experiment, only the setup is a bit more difficult and unfortunately so are the updates.

4.2 Alternative experiments

The biggest problems I expect for this experiment are the two updates it requires. I can explain this best by using an example.

If you look at the evening news at two different channels, you will see two different weathermen. Now, if both of the weathermen predict rain, it is quite easy to conclude it is probably going to rain. The model, however, does not require the subjects to draw a conclusion about the rain. It requires them to conclude that both weathermen send the same message and thus it is more likely that they are capable then it was before you noticed this. At first sight this does not appear immediately clear, but if one thinks of it in this way it will: if they disagree, one of them has to be wrong, if they agree however, both of them can be right.

After that, somewhat hard, but straightforward analysis, the model requires even more of the experts. It adds an extra layer of rationality. Where the public has to conclude on something that is left unspoken, the weathermen (experts) are actually expected to alter their message about the weather in some occasions to make the public believe they are able weathermen. The model requires the weathermen to coordinate.

These updates and especially the extra layer of rationality, tested apart from the rest of the model, could make interesting experiments on their own.
For the following experiments I will constantly use the weathermen example. This is by far not the only situation possible in the experiments, other examples are given in Appendix B.

4.2.1 Rationality of the public

Class room

In the main experiment the public has to conclude how valuable a member in a group of experts is based on his actions, his statements, and the actions and statements of the rest of the experts. Much of the individual decision literature already focuses on subjects (dis-)ability to do a proper Bayesian update. Since the main experiment poses a new problem and new layer of rationality some (re-)testing can be done. Presenting a class room of students with similar problems and asking them what they think and how they approach this update might provide us with some inside information about how people (and subjects) approach this problem.
The questions posed could be as simple as:

On some Monday the weathermen on channel 1 and 3 agreed that it is going to rain that Wednesday. The weatherman on channel 2 however said it was going to stay dry. That Tuesday you are packing for a trip, you are leaving on Thursday and can only bring one set of clothing. The weatherman on channel 1 predicts sunshine, the one on channel 2 however predicts rain. You cannot check the prediction of the third weatherman, but you absolutely want to pack the clothing that fits best with the weather. For what kind of weather should you pack?

If the subjects use the available information, signal concurrence on the part of the weathermen on channel 1 and 3, they should conclude that it is most likely that the channel 1 weatherman is good and pack for a dry day. If, however, they want to “play it safe”, they will probably pack for rain. This is a simple example, and in the class-room it could be used to explain updating of information.
The two biggest drawbacks here are, first, it is all cheap talk, second, the model behind this example is implicit.
The first problem could be solved by betting, see below. Solving the second problem would require adding more information. The problem would then start to look like:

In some time of the year it rains about 30% of the days. A Monday during that period the weathermen on channel 1 and 3 agreed that it is going to rain that Wednesday. The weatherman on channel 2 however, sais it is going to stay dry. That Tuesday you are packing for a trip you have planned for Thursday and you can only bring one set of clothing. The weatherman on channel 1 predicts sunshine, the one on channel 2, however, predicts rain. You cannot check the prediction of the third weatherman, but you absolutely want to pack the clothing that fits best with the weather. What should you pack clothing for, a rainy day of for a sunny day?

Adding this information does come at a cost, it costs more time to explain and the update gets more complex.

Betting

The setup of the experiment above is simple, but revolves around cheap talk from the subjects. The experiment could easily be transformed in an incentive compatible experiment by having the subject bet on the outcomes, to make them put their money where their mouth is. To stay with the simple weathermen example a bet could look like this:

There are three weather services in country X, for convenience we will call them K, L and M. Three days a go K and L predicted that it would rain the next day, M said it would remain dry. Two days a go M said it would rain the next day (yesterday), K said if would remain dry.

Given this information what is more likely:
A) Get 10 ECU if it rain was more likely
B) Get 10 ECU if dry weather was more likely

If the subjects pick up on the signal concurrence again, they should pick option B.
Again, adding more information makes the model more complete. Saying it rains 30% of the days, half of the weathermen make it up as they go, etc., could complete the picture, but at the expense of a more complex update.

4.2.2 Rationality of the experts

It is quite hard to get the experts to persuade a non-existing public into believing they are able. Therefore, the experimenter will have to either take the role of the public, or postulate the behavior of the public so the subjects in their role of experts can predict the possible reactions of the public and act accordingly.

Computerized public

A viable way of modeling the public could be to have a computer player take that role. The experiment would pitch the subjects against a computer player that can calculate the true expected value of his actions. That way the subjects can, if they understand the model, predict what their opponent will do and are able to use this to their advantage.

If the same experiment and information is used as in the main-experiment's reputation treatment, it would be possible to see if there is different behavior between the experts in the human-monitor setup compared to a setup with a rational computer-monitor setup.

The hypothesis that can be tested in this way, is whether or not the subjects are capable of the higher level of rationality.

- If, there is no difference between the computer-monitor group and the human-monitor group and the theory of Visser & Swank (2007) holds, than we would know that the theory is correct and subjects are capable of the higher level of rationality and anticipate rational behavior of their opponents to some degree.
- If, the theory holds for the reputation with computer-monitor treatment, but does not hold, or does not hold as strongly, for the reputation with human-monitor treatment, than we know that the experts take the limited rationality of human monitors into account and are capable of the higher level of rationality to some degree.
- If, the theory does not hold in any setup, than we might have to conclude that we cannot get subjects to reach this extra layer of rationality in this setup and need to find out why.

A possible weakness to this setup is the definition of rational behavior for the computer-monitor. The easiest approach would be to have the computer play the Perfect Bayesian Equilibrium and see if the subjects can find it. The problem with this approach is that the computer would not really be playing rational, and the subjects' rational reaction to this would be a pure strategy.

A second approach is having the computer perform Bayesian updates to determine his beliefs, which is as close to truly rational as possible. The problem is then to determine the precise update the computer should do.

If it uses the data of former actions by the subjects, the one-shot nature of the game is lost, which is a big deviation from the model.

If the computer uses all available data from the round (the actions of all subjects) he might be comparing people that use a totally different strategy. That way the test would mostly be about rational behavior in the group instead of individual rationality. Luckily this rationality within the group is a strong indication of individual rationality, since the individuals that make up the group apparently act in a way that resembles rationality.

Personally I think this last option is the best. This way the computer uses the observed behavior of the group to determine his beliefs or priors, and thus his priors are averages, which is fair for a prior. It also means that if all members of the group play the Perfect Bayesian Equilibrium, they will force the computer to do so as well. On the other hand, if the group of subjects play a different strategy, so will the computer. Due to the repetitive nature of the experiment a lot of rounds will be created. This might cause the subject to apparently deviate from equilibrium in some rounds, due to the mixed strategy the subjects should play. That is, in some rounds too much or not enough bluff will be found, simply due to random effects, not due to real deviations. This implies that if we can find the rational behavior in this setup, the rational behavior is quite robust.

5. Informal testing of the model

Since setting up the lab will take some time and lab-time and subjects should not be wasted, I have tried setting up a variation to the main experiment in a way that does not require a lab. These test-runs might not deliver the data I need for analysis, they do provide some insight in the behavior of people confronted with the same choices as the subjects have to make in the lab. To ensure I do not use the same subjects twice, these test-runs have been done with people that do not follow any courses at the Erasmus University, so cannot enroll in the university's paid subjects pool. In this experiment I tried to stay as close to the main experiment as possible, there is however one major deviation. In the main experiment all subjects play only one part, however, in this game, to keep it fair, players change roles each round.

5.1 Preparation

I took 4 decks of cards that have identical face-side pictures, but different color backs, half blue, half red which indicated the blue and the red state. To each of twelve rounds, that is each of the sets of 2's 3's … K's, I randomly assigned states and signal qualities.
Blue state/card backs meant investment would be beneficial, red meant that investment would fail.
The two signal distributions for the two committee members were separated by types of the face side of the cards; hearts and diamonds made up one signal distribution, clubs and spade the other. The left-hand side of the table below shows part of the table I used when I first setup the game, the right-hand side is the same table only now translated to the cards used in the game.

	Random draws			Cards		
Round	State	Expert 1	Expert 2	State	Diamonds/Hearts state-set	Clubs/Spades state-set
2	1	1	0	Blue	4*blue	2*blue, 2*red
3	0	0	1	Red	2*blue, 2*red	4*red
4	0	1	1	Red	4*red	4*red

The random draws were made in such a fashion that they created a 50/50 dispersion of good and bad analysts, combined with a 50/50 distribution of the states.

After putting together these combinations for each round (2,3….K), except aces, and adding some monopoly money to use as stakes, the game can be played.

5.2 The game, single treatment

Three players compete in this game. In each round two of them are a committee and the third one is bank. The different actions in the game are listed, in order, below.

1. The bank takes one of the rounds (2…K), containing two sets of four cards, the state-sets. The bank hands the four clubs/spades to one committee member and the four diamonds/hearts to the other. Both committee members draw a card from his/her state-set. The back of the card proxies for their private signal and is either blue or red.

2. In the deliberation state the committee members use the ace cards to send each other a message about the state they think they are in. Blue (red) card-back means blue (red) state, and the cards are returned. The exchange is done in such a fashion that the bank cannot see the back of the cards (the color).

3. Since voting is private information, the two committee members vote with the ace cards. Both add one covered ace card to a vote stack, then they check it. If both cards are blue, they open up the vote stack and the investment is made by having both committee members pay a stake to the experimenter, otherwise the vote stack remains closed.

4. Both committee members then send a message to the bank by giving him an ace card, indicating their individual statements about the state. Again blue cards are used to claim a blue state and red card to claim a red state.

5. The bank can choose to buy the state-set of one committee member, no committee member, or both committee members. If the bank buys a state-set, he gives twice the stake to the owner of that set. Each committee member that the bank buys from, turns his/her state-set over. If the state-set contains 4 cards of the same color (blue/red), then the pot pays 4 times the stake to the bank. If the distribution is 2/2 no payment to the bank is made. The committee members keep what is paid to them in all situations.

6. If investment was made by the committee, any unturned state-sets are also turned up and if 6 or 8 of the state cards are blue, the investment stake is returned to the committees plus 50% from the experimenter. If the distribution is 4/4 the committee members get their stake back. Else the stake is lost to the pot.

7. All cards are returned and the state cards are replaced by a new round. Roles are switched.

Experience

Playing this game twice, both with people with college grade or academic backgrounds, yielded some observations:

1. The speaking with one voice comes quite naturally in this setup, it is clear that the committee members want the bank to buy their state-sets.
2. It took a couple of rounds before one of the other players realized this and noticed that there was no information in the message to the bank, but there should be some information in the investment decision.
3. Some rounds (and probably too much information from my side) later, the same player noticed that a positive investment decision from the committee was a positive signal for the bank to buy state-sets. Since his committee member was a bit more conservative (only invest if two positive signals were received) he decided to bluff and sent a positive message even though he drew a negative signal.
4. The game has a limited amount of options and is therefore quite dull and gets weary quite quickly.
5. There was a clear order in which messages from different players were trusted. The one pulling of the bluff was known to do so in other games as well, so his state-set was bought les frequently than the others.

Playing this game face to face stimulates the searching for tricks to outwit each other. It also shows that the speaking in one voice prediction follows directly from the proposed setup of the model. The trick mentioned here, the bluff, is actually an effect the theory only predicts if there is a different valuation of the reputation (sale of the state-set), but is found here due to interaction between the players.

The game proved two things. First if the players think about the game and are allowed to reason about it, they might be able to see trough the setup and the information structure and find the equilibrium, al be it at great difficulty. Second, testing it face to face does add some interaction that can facilitate the search for the necessary 'tricks', it also creates some unpredictable noise, since players have/form expectations toward each others' behavior that are entirely unrelated to the experiment.

5.3 The game, mixed treatments

The game above is basically a single treatment game and it gets a bit dull, I felt adding more treatments was a way to make it more interesting. Rolling a dice at the beginning of the round to determine the price of the state-sets seemed to be an improvement, since it created different treatments and more variability.

By setting the prices in such a way that in some of the rounds the payoff for the committee of selling the state-sets is lower than the payoff of investing and vice versa, I expected to find different strategies being played.

This was achieved by setting a fixed investment stake of 100, and setting the price of the state-sets between 50-300 (50*d6). I was hoping to see more variability in the strategies played this way, I found something else.

Experience

Playing the mixed treatment game with new subjects (5 new subjects, two games) proved to be slightly more difficult then before. For one, it took me longer to explain the game to them now I was watching the amount of information I gave them about the information structure and its' implications. Also, the amount of conclusions I am willing to draw from this experience is a lot smaller, since this group seemed unable to determine a working strategy for themselves. Subjects in this round were college grade and below. The observations from these rounds:

1. Although most of the players matched their message to the bank to their investment decision, not all players did so (one did not), leading to a spectacular breakdown of the speaking with one voice outcome. Most peculiar was, that it did not seem to negatively affect the rate with which state-sets were bought from him (he actually sold most state-sets). While talking about the game, players demonstrated they could grasp the fact that the message to the bank contained no information, this could explain the non-equilibrium message strategy, it did not seem to matter.
2. Non of the subjects picked up on investment as a sign of signal concurrence until this was explained to them. The difference in sell rates of the state-sets is non-significant

(5 sets were sold after investment, 4 without investment).

3. The 50/50 chance distribution might make it easier to calculate the expected values, it also creates a situation were out-of-equilibrium behavior is not punished harshly, there usually is a substantial chance it goes either way.

Adding the dice did not improve the game, it just provided some more clutter to the thinking the subjects have to do to be able to determine a working strategy. This made it clear that any information that will be provided to the subjects in the real experiments should be precise and complete, if we want them to understand the experiment correctly.

The observation on the 50/50 chance distributions is troublesome, changing the chance distribution seems like making the calculations more difficult and thus making investment and buying analysts more scary (ambiguity aversion), but it might be necessary if we want the subjects to learn from mistakes quickly. Again clear information to the subjects might be the best (partial) solution to this problem. This also entails only giving one treatment to each subject, so he can learn the subtleties of that particular treatment quicker. I do believe proper incentives can help this experiment along, mucking about with messages, since it does not really matter anyway, should be less appealing ones real money is at stake. This should make some of the effects found easier to identify.

6. Concluding remarks

Finding a way to test the theory of Visser & Swank (2007) proved to be quite a challenge. First digging in to related literature and the experimental literature both in economics and social psychology and then trying to apply what I've learned opened up a new scope of research for me.

After visiting the University of Amsterdam, where Randolph Sloof, Silvia Dominguez Martinez and Ferdinand von Siemens were kind enough to let us observe one of their own experiments, I realize that before this research is finished a lot of work will have to be done. In setting the experiment up for the computers in the lab I might find some things that need to be refined, the data collection will take quite some time and then the processing begins. Now I have been thinking about the actual experiment for quite sometime, however, it will be quite rewarding to actually do the experiment.

Before that step is taken, it will probably be worthwhile to do some test runs in the form of the other experiments, just to get a better feeling avout what it is all about and where the problems lie, especially since this will be the first experiment I ever conduct.

About the outcomes of the tests of the theory itself I've got mixed feelings, especially since the informal tests I did. The updates it requires are quite hard, even more so, if you have no knowledge of the theory behind the updates, or even have never had any formal training in updating probabilities. I do believe it is likely that parts of the outcome appear in the data, especially the speaking in one voice outcome seems logical, but I have sincere doubts about the ability the average person has in reaching the higher level of rationality needed to anticipate rational updates of another player. For me that means, that the outcomes of the alternative experiments in that area might really become interesting. Finding setups where the subjects can see trough the setup in such a way that they are capable of this update might prove challenging though.

References

1. Arkes & Ayton, 1999, *The Sunk Cost and Concorde Effects: Are Humans Less Rational Than Lower Animals?,* Psychological Bulletin, vol. 125 no. 5, pp. 591-600.
2. Baron & Byrne, 1981, *Social psychology: Understanding human interaction*, Boston, Allyn & Bacon
3. Bazerman , Giuliano, Appelman,1984, *Escalation of commitment in individual and group decision making*, Organizational behavior and human performance, vol. 33, no. 2, pp. 141-152
4. Bobocel & Meyer, 1994, *Escalating Commitment to a Failing Course of Action: Separating the Roles of Choice and Justification* Journal of Applied Psychology, vol. 79 no. 3, pp. 360-363
5. Brockner, 1992, *The escalation of commitment towards a failing course of action: Towards theoretical progress,* Academy of Management Review, vol. 17, no.1, pp. 9-61.
6. Brockner & Rubin, 1985, *Entrapment in escalating conflicts, A Social psychological Analysis*, SSSP
7. Brockner, Rubin & Lang, 1981, *Face-saving and entrapment,* Journal of Experimental social Psychology, vol. 17, pp. 247-266
8. Gerling, Kertin, Grüner, Kiel, Shulte, 2005, "*Information acquisition and decision making in committees: a Survey",* European Journal of Political Economy, vol. XVI, pp 563-597
9. Harrison and Harrell, 1993, *Impact of "adverse selection" on managers' project evaluation decisions,* Academy of Management Journal, vol. 36 no. 3 pp. 635-43
10. Janis, 1972, *Victims of groupthink: a Psychological Study of Foreign-Policy Decisions and Fiascos*, Boston: Houghton Mifflin Company
11. Kahneman & Tversky, 1972, *Subjective probability: A judgment of representativeness*, Cognitive Psychology vol. 3, pp. 430-454.
12. Kahneman, & Tversky, 1973, *On the psychology of prediction*, Psychological Review vol. 80, pp. 237-251.

13. Kahneman & Tversky, 1979, *Prospect theory: An analysis of decision under risk*, Econometrica vol. 47, pp. 263-29

14. Kanodia, Bushman and Dickhaut, 1989, *Escalation Errors and the Sunk Cost Effect: An Explanation Based on Reputation and Information Asymmetries*, Journal of Accounting Research vol. 27 No. 1. pp 59-77

15. Moon, Conlon, Humphrey, Quigley, Devers and Nowakowski¸ 2003, *Group decision process and incrementalism in organizational decision making,* Organizational Behavior and Human Decision Processes, vol. 92, pp. 67–79

16. Sharp and Salter, 1992, *Project Escalation and Sunk Costs: A Test of the International Generalizability of Agency and Prospect Theories,* Journal of International Business Studies, vol. 28, no. 1, pp. 101-121

17. Staw, 1976, *Knee-deep in the big mud: A study of escalating commitment to a chosen course of action*, Organizational Behavior and Human Performance, vol. 16, pp. 27-44

18. Staw & Fox, 1977, *Escalation: The Determinants of Commitment to a Chosen Course of Action*, Human Relations, vol. 30, pp. 431-450.

19. Swank, Swank & Visser, 2008, *How Committees of Experts Interact with the Outside World: Some Theory, and Evidence from the FOMC*, Journal of the European Economic Association, MIT Press, vol. 6, pages 478-486

20. Tan & Yates , 1995, *Sunk cost effects: The influences of instruction and future return estimates,* Organizational behavior and human decision processes, vol. 63, no 3, pp 311-319

21. Teger, 1980, *Too much invested to quit,* New York: Pergamon Press

22. Visser & Swank, 2007, *On Committees of Experts,* Quarterly journal of economics, vol. 122, pp. 337-372

23. Wong and Kwong, 2007, *The role of anticipated regret in escalation of commitment,* Journal of applied psychology, vol. 92, no. 2, pp 545-554

24. Whyte, 1993, *Escalating commitment in individual and group decision making: a prospect theory approach*, Organizational Behavior and Human Decision Processes, vol. 54, pp. 430–455.

25. Wilson and Zhang, 1997, *Entrapment and escalating commitment in investment decision making: A review,* British Accounting Review, vol. 29, pp. 277–305

Appendix A

Updating the prior, state

As said in the main text there are two types of situations, the members received identical signals, or they received different signals.

Signal concurrence can occur in two ways, they either both received the correct signal, or they both got it wrong. In the setup in the experiment the chance that an expert receives the correct signal is 0.75 (π+0.5(1- π)) and the chance an expert receives an incorrect signal is therefore 0.25. A Bayesian update then tells us that if both experts receive the signal investing is good (bad), there is a 0.9 chance the investment will succeed (fail).

In case of mixed signals the update is even simpler. Since both experts are ex ante equally likely to be competent, their signals cancel each other out. The posterior is equal to the prior, which is 50%.

Updating the prior, ability

Recapping from Visser & Swank (2007), $\hat{\pi}(X = 1; \beta) = \Pr\,(sm \mid X = 1; \beta)$, with $\hat{\pi}$ being the update probability a member is smart (sm), X=1 denoting a positive investment decision and β the probability the committee invests despite mixed signals.

Calculating the poserior using Bayes' rule requires Pr(X=1), Pr(X=1|sm) and Pr(sm) = π.

Note that:

$$\Pr(X = 1|t_1 = sm, \mu = u) = \frac{1+\pi}{2} + \frac{1-\pi}{2}\beta$$

$$\Pr(X = 1|t_1 = sm, \mu = -u) = \frac{1-\pi}{2}\beta$$

$$\Pr(X = 1|t_1 = du, \mu = u) = \frac{1+\pi}{2}\left(\frac{1}{2} + \frac{1}{2}\beta\right) + \frac{1-\pi}{2}\frac{1}{2}\beta$$

$$\Pr(X = 1|t_1 = du, \mu = -u) = \frac{1-\pi}{2}\left(\frac{1}{2} + \frac{1}{2}\beta\right) + \frac{1+\pi}{2}\frac{1}{2}\beta$$

Given that all of these situations occur with probability 0.25 this adds up to

$$\Pr(X = 1 | t_1 = sm) = \frac{1}{4}(1 + \pi + 2(1 - \pi)\beta)$$

$$\Pr(X = 1 | t_1 = dm) = \frac{1}{4}(1 + 2\beta)$$

And then

$$\Pr(X = 1) = \frac{1}{4}(1 + \pi + 2(1 - \pi)\beta)\pi + (1 - \pi)\frac{1}{4}(1 + 2\beta)$$

$$\Pr(X = 1) = \frac{1}{4}[\pi + \pi^2 + 2\pi\beta - 2\beta\pi^2 + 1 + 2\beta - \pi - 2\beta\pi]$$

$$\Pr(X = 1) = \frac{1}{4}[1 + \pi^2 + 2(1 - \pi^2)\beta]$$

So using Bayes' rule

$$\hat{\pi}(X = 1; \beta) = \Pr(sm | X = 1; \beta) = \frac{\frac{1}{4}(1 + \pi + 2(1 - \pi)\beta)}{\frac{1}{4}[1 + \pi^2 + 2(1 - \pi^2)\beta]}\pi = \frac{1 + \pi + 2(1 - \pi)\beta}{1 + \pi^2 + 2(1 - \pi^2)\beta}\pi$$

Using Pr(X=0)=1-Pr(X=1)

$$\hat{\pi}(X = 0; \beta) = \Pr(sm | X = 0; \beta) = \frac{1 - \frac{1}{4}(1 + \pi + 2(1 - \pi)\beta)}{1 - \frac{1}{4}[1 + \pi^2 + 2(1 - \pi^2)\beta]}\pi = \frac{3 - \pi - 2(1 - \pi)\beta}{3 - \pi^2 - 2(1 - \pi^2)\beta}\pi$$

Appendix B

Since there are many situations in real life where we trust on someone's advice on the grounds that he knows more then we do, so there could also be many scenarios that can be used for these experiments. Below is a list with a couple of examples, it is not complete and I only give the version that fits with the simple class room setup from 4.2.1, but it could provide a good starting point.

Restaurant guides

Each year several restaurant guides publish ratings of the restaurants in your neighbourhood. Given that you have a big date coming up you decide to consult a couple of those guides. There are several establishments you are looking at, two Italian restaurants an African restaurant and the Japanese restaurant located next to it. You decide to consult the three biggest independent guides you can find online. All of them have ratings on the Italian and Japanese restaurants, but only two of them list the African restaurant. This is what you find

	Guide 1	Guide 2	Guide 3
Italian I	Good	Good	Mediocre
Itallian II	Mediocre	Mediocre	Good
African		Good	Mediocre
Japanese		Mediocre	Good

After some carefully asked questions you find out that your date thinks going to an Italian restaurant is a bit cheesy, but that your date seems equally interested in the African and the Japanese restaurant, as long it is a good restaurant.

Assume you do not have a particular preference for either one and both are within budget and reach, which restaurant would you judge most likely to be good?

There is a clear match in signals between 1 and 2 so it is most likely the African restaurant is the better one.

Investment banking

As an entrepreneurial student that needs slightly more money than the study grants the government provides, you decide to use what is left of your saving account for short term investing to create some extra income. Since there are plenty of stock advisors, you decide to consult some of them. The stock market, however, requires quick reactions to make an actual profit, so you restrict searches to 4 websites you can check quickly. The stocks they pick as todays definite winners are (in no particular order)

Site A	**Site B**	**Site C**	**Site D**
Non-existent ltd	Fictitious inc.	Non-existent ltd	Fairy-tale corp.
Fictitious inc.	Non-existent ltd		Fictitious inc

As <u>losers</u> are listed

Site A	**Site B**	**Site C**	**Site D**
Belly-up investments	Fairy-tale corp.	Illusory corp.	Non-existent ltd
Imaginary ltd.	Imaginary ltd.		Illusory corp.

Before the market opens, Fairy-tale corp. announced a press meeting later that day. Your gut tells you this is the one for today, but will you invest in it, or is it going to drop, in which case you can better short sell it. Non of the sites change the listing of Fairy-tale corp.

Short selling is in line with the predictions of Site A, Site A agrees with B and C in its listing, while D shows signal incongruence in case of Non-existent ltd. So short selling is most likely the best option given the information in the setup.

Betting on sports

Your co-workers have a lively soccer-pool for the championship, and since they seem to be enjoying it, you are thinking about joining. They are also already betting for the game coming Sunday, which could be a fun to start with, since you are too late to join the pool for the championship anyhow. Your office mates seem to be in a bit of a predicament, two of them are predicting that the team D is going to win the championship, while a third one is certain that it will be Team H. when the debate is settled down a bit it is your turn to place a bet for the game that Sunday between L & R. You want to win the bet, but you don't know the clubs that are playing. Two the colleagues that were just in the hefty debate come to your rescue with some advice, of course they have both studied the strategies of both teams and have a thorough analysis. The Team H supporter is therefore certain that Team L is going to win next Sunday, while the other one (one of the Team D supporters) is putting his money on Team R.

Signal concurrence would suggest that Team R stands a better chance.