Collective Delusion in Risk Management

Augustin Landier[*] David Sraer[†] David Thesmar[‡]

December 26, 2008

Abstract

We model an organization as a two-agent hierarchy: a Decision Maker in charge of selecting projects and an Implementer in charge of their execution. Both have intrinsic preferences over projects and have the ability to manipulate their beliefs on the probability of success of these projects. We find that the presence of endogenous beliefs reduces the ability of the organization to use outside, objective information about projects' quality. It also gives rise to multiple equilibria: a well functioning chain of command may turn into a collectively delusional hierarchy unable to process external information. We apply these insights to explain risk management failures in the context of the recent financial crisis.

[*]NYU-Stern, alandier@stern.nyu.edu

[†]UC Berkeley, sraer@berkeley.edu

[‡]HEC, CEPR and ECGI. thesmar@hec.fr

1 Introduction

The recent turmoil on credit markets has drawn attention to the risk management function. On many trading floors around the world, traders have been writing insurance against rare events: examples include keeping long positions on CDO tranches or selling protection against default (CDS). In normal times, it is the role of risk management to ensure that the received insurance premia are not entirely considered as income, but that enough capital is set aside to protect the institution against the risk that it is taking. In the period that led to the current crisis, however, risk management has failed to play this role.[1]

One possible reason is that risk managers did not have enough support from their hierarchy. Traders were seen as profit centers, while risk management only generated costs in the short run.[2] This begs the question, however, of why top management, who sometimes had most of its own wealth invested in the bank, accepted to take on so much risks. René Stulz (2008) recalls that managers of Bear Sterns and Lehman owned millions of dollars of equity, and partners at LTCM collectively had $2bn invested in the fund before it collapsed.

An alternative hypothesis, which we explore in the present paper, is that risk managers and traders simply became genuinely over-optimistic about the risks credit derivatives were exposing their firm to. For instance, the *UBS shareholder report* (2008) suggests that "Fundamental analysis of the subprime market seems to have been generally based on the business' view and less on Market Risk Control's independent assessment. In particular, there is no

[1]For instance, the *UBS shareholder report on write downs* (2008) issued in April 2008 recalls that the IB fixed income division had no senior risk management officer in 2006. In December 2007, its CDO desk ended up with a loss of about $12bn in 2007 through various long exposures to subprime mortgages.

[2]Ibid.

indication that MRC was seeking views from other sources than business." In other words, risk managers took risk evaluations emanating from the CDO desk at face value, without even questioning the potential conflicts of interests. Yet, risk managers were in principle independent, since neither their compensation nor career path were tied to the desk P&L.

We propose a theory explaining how and why such collective blindness may emerge. This theory rests on two building blocks. First, following Augustin Landier et al (2009), we model the Trader - Risk Manager relation as one between a Decision Maker (Trader) selecting a project (an asset) and an Implementer (Risk Manager) in charge of its execution (determining the scale of the Trader's investment). Both the Trader and the Risk Manager have intrinsic preferences for one of the assets. As in Landier et al., a Risk Manager with dissenting intrinsic preferences emulate more efficient project selection (i.e. asset purchases) by compelling the Trader to pay more attention to objective information over intrinsic preferences. Second, we allow both agents to optimally choose their interpretation of signals about asset returns. We use the approach of "anticipatory utility" (Augustin Landier, 2000, Roland Bénabou and Jean Tirole, 2002, Markus Brunnermeier and Jonathan Parker, 2005), which embeds the following tradeoff: on the one hand, forging one's memory can increase utility by letting the agent "savor" a specific outcome, but on the other hand it induces poor decision making.

Collective delusion limits the extent to which signals are used in the decision making process. This result operates through a positive feedback loop between incentives and wishful thinking. Assume the Trader is intrinsically biased toward asset 1 (e.g., the risky asset). The Risk Manager anticipates that she will select this asset more often. Because the Risk Manager cannot change the portfolio composition, he might as well try to savor the prospect of asset 1 by forgetting any bad information on this asset. The more important is "savoring" in the

Risk Manager's utility, the more he will expand the size of the investment in asset 1, which in turn will tilt the Trader's incentives toward picking the risky asset. Such complementarities naturally generate multiple equilibria: a "reactive" chain of command – one that always follows outside information – may turn into a collectively delusional hierarchy that is not able to process external information anymore.

Our paper deviates from the finance literature on risk management by shifting the focus of the analysis from methods of risk evaluation (see for instance Anthony Saunders and Marcia Cornett, 2006), towards organization design. We believe this is an important contribution as recent risk management failures seem to have been driven as much by organizational issues as by the inherent difficulty to measure credit risk. For instance, citing the UBS report, René Stulz (2008) argues that communication of risk exposure to the top management was highly inefficient.

This paper is also part of the nascent literature on behavioral organization. It is closely related to Roland Bénabou (2008); like us, he presents the model of an organization with several agents who can engage in reality denial. The organization he focuses on is very different from ours: in his model efforts are simultaneous and additive, while we focus on an asymmetric chain of command where the Decision Maker and the Implementer act sequentially. Like us however, he finds that under certain conditions, delusion strategies between members of one organization are strategic complement. While in our model such complementarity is naturally built in the hierarchical structure, in his model it arises even in the absence of complementarity in effort. Another related paper is Eric Van Den Steen (2007), who shows that divergence in opinions in organizations may make high powered incentives counterproductive. In his paper, however, beliefs are exogenously fixed, while we endogenize reality denial.

This paper has three remaining Sections. Section 2 presents the basic set-up. Section 3 defines and characterizes the equilibria. Sections 4 gives various comparative statics results.

2 The Basic Model

2.1 Set-Up

The basic set-up follows Landier et al. (2009). The organization consists of two risk neutral agents: an Implementer (he) and a Decision Maker (she). There are two states of Nature S labelled 1 or 2, both occuring with probability 1/2. The Decision Maker selects one of two projects labeled 1 or 2. The Implementer is in charge of implementing it, and can put in high, or low, effort. The cost effort $\widetilde{c}$ is borne by him and is drawn from a distribution of c.d.f. F, assumed to be twice differentiable and concave. The probability of success of a given project depends both on the Implementer's effort and the state of nature: if the project coincides with the state of Nature, and the Implementer puts in high effort, the project succeeds. In this case, it brings R to the owner of the organization. Profit is $-L$ when the project fails. In our risk-management interpretation, the Implementer is the risk-manager and the cost $\widetilde{c}$ can be interpreted as the cost of capital for the project.

The tension in this model comes from the fact that the Decision Maker has intrinsic preferences for project 1, while the Implementer may have different intrinsic preferences. Neither the Decision Maker, nor the Implementer, receive any payoff when the project fails, e.g. because of a Limited Liability rule. Therefore, when project j succeeds, the Decision

Maker receives utility B_j, and the Implementer b_j, where $B_1 > B_2$. Furthermore, in order to simplify the exposition, we assume that the Decision Maker is "more biased" than the Implementer in the following sense:

$$F\left(\frac{b_1}{2}\right) B_1 > F\left(b_2\right) B_2 \tag{1}$$

This will ensures that when the signal is 1, the Decision Maker always selects project 1, even if the Implementer intrinsically dislikes it. It allows us to focus on reactivity to signal 2 only.s

2.2 Timing

The sequence of actions has five different steps. At date 1, there is a publicly available signal σ on the state of nature. The probability that the signal reveals the true state of Nature is $P(\sigma = i|S = i) = \alpha > 1/2$. At date 2, both members of the organization decide whether they alter their memories or not (as in Bénabou and Tirole, 2002; see below): each one of them chooses which signal $m \in \{1, 2\}$ he or she will remember. At date 3, relying on her memory of the signal, the DM chooses which project to order. In period 4, the Implementer learns his cost of high effort $\widetilde{c}$. Relying on his memory of the signal, he decides on the effort to exert. In the final period 5, payoffs are realized.

2.3 Wishful Thinking and Information Structure

In Landier et al (2009), both the DM and Implementer hold realistic beliefs about the probability of success of the project. These beliefs are formed after observation of the informative public signal. In the present paper, we allow agents to distort their beliefs using the memory altering technology first developped in Bénabou and Tirole (2002). Once each agent has observed the signal, he or she can choose to completely change her memory thereof. Once in period 3, the Decision Maker will select the project that provides her with the largest expected utility, based on her (potentially false) memory of the signal. In period 4, the Implementer's choice of effort will similarly be made using his own (potentially) false memory of the signal.

As in Bénabou and Tirole (2002), Brunnermeier and Parker (2005), or Landier (2000), agents select the beliefs that maximize the sum of two expected utilities: (1) the expected utility using the forged, ex post belief [3] (2) the expected utility using the accurate, ex ante beliefs. The second term in the utility represents a realistic assessment of the distortions that subjective beliefs will create. s is the weight of the savoring term in the utility, and we normalize to 1 the weight of the realistic utility. To clarify exposition, we allow s to differ for the Implementer s_I from the Decision Maker s_{DM}.

We then model the memory recollection process. To clarify exposition, we assume that agents are fully *naive*: in periods 3 and 4, agents fully trust their memory and are thus naive about the way memories are forged.[4] To further simplify exposition of the main effects, we allow for "fuzzy memory" strategies: period 2 agents can choose the probability that their future selves will place in the signal being 1. This probability can be set anywhere between zero and one.Finally, we make the natural assumption that memories are private

[3] This is the utility that will actually be experienced by the agent once his/her memory will be altered. The assumption here is that rational agent can "savor" such a utility as of period 2

[4] This is in contrast with Bénabou and Tirole (2002) who assume full sophistication: in their paper, agents expect their memory to have been altered, and update accordingly. None of our result hinges on our full naivete assumption

information.

A feature of our model is that having a risk-averse Risk-manager does not guaranteee that excessivly risky projects are avoided as the beliefs of the trader contaminate the beliefs of the risk-manager. Such contagion of beliefs is endogenized through the agent's anticipation utility, which gives them a preference to distort memories to feel better about the future.

3 Equilibria

3.1 Equilibrium Definition

The equilibrium concept is subgame perfect Nash equilibrium. For the Implementer, the strategy space is made of two functions:

$$\mathcal{R}_I \; : \; \sigma \in \{1;2\} \mapsto \rho_I \in \{1;2\}$$

$$E \; : \; (o, \rho_I) \in \{1;2\} \times \{1;2\} \mapsto e \in \{\text{High}; \text{Low}\}$$

where $\mathcal{R}_I$ is the implementer's belief choice, once the signal is observed: at date 4, he will believe that signal was 1 with probability ρ_I. The Implementer chooses this probability at $t = 2$. E represents the Implementer's level of effort, which depends on his memory choice and on the project selected by the Decision Maker.

The strategy space of the Decision Maker is given by:

$$\mathcal{R}_{DM} \; : \; \sigma \in \{1;2\} \mapsto \rho_{DM} \in \{1;2\}$$

$$O \; : \; (\rho_{DM}, e) \in \{1;2\} \times \{\text{High}; \text{Low}\} \mapsto o \in \{1;2\}$$

where $\mathcal{R}_{DM}$ is the Decision Maker's belief choice, once the signal is observed: at date 3, she will believe that the signal was 1 with probability ρ_{DM}. O is the project she selects, which depends on her own recollection of the signal, as well as on her expectation of the Implementer's level of effort. We next define the equilibrium:

Definition 1 *An equilibrium is a set of functions* $(\mathcal{R}_{DM}, \mathcal{R}_I, O, E)$*, such that:*

1. E maximizes the expected utility of the Implementer, using ex post beliefs defined by $\mathcal{R}_I$.

2. O maximizes the Decision Maker's expected utility, using ex post beliefs defined by $\mathcal{R}_{DM}$

3. $\mathcal{R}_I$ maximizes the sum of s_I times the expected utility using ex post beliefs from $\mathcal{R}_I$ plus the expected utility using objective, ex ante beliefs.

4. $\mathcal{R}_{DM}$ maximizes the sum of s_{DM} times the expected utility using ex post beliefs from $\mathcal{R}_{DM}$ plus the expected utility using objective beliefs.

3.2 Reactive Equilibria

We first characterize reactive equilibria, i.e. organizations where the Decision Maker always selects the project most likely to succeed (project 1 with signal 1, project 2 with signal 2).

Proposition 2 *Reactivity and Realism*

An equilibrium with full reactivity exists if and only if

$$(1 + s_{DM})\, \alpha F(\alpha b_2) B_2 \geq ((1-\alpha)F((1-\alpha)b_1) + s_{DM}\alpha F(\alpha b_1))\, B_1 \qquad (2)$$

In this case, both the Implementer and the Decision Maker hold realistic beliefs.

Proof. See Appendix 7.1. ■

The above condition ensures that, when faced with signal 2, a realistic Decision Maker prefers to select project 2 over project 1.[5] Assume the signal is 2. The left hand side of the inequality is her realistic payoff of ordering 2: in this case, the objective and subjective utilities are identical, project 2 has a probability α of success, the (realistic) implementer will thus be implementing it with probability $F(\alpha b_2)$, and the final payoff in case of success to the DM is B_2. The right hand side of the inequality represents the Decision Maker's expected payoff of selecting project 1 and remembering signal 1. The first term is the realistic expected utility: $(1-\alpha)F((1-\alpha)b_1)B_1$, which accounts for the fact that project 2, which has a probability $1-\alpha$ to succeed, will only be implemented with probability $F((1-\alpha)b_1)$ and generate the high payoff B_1. The second term is the anticipatory utility: $\alpha F(\alpha b_1)B_1$ multiplied by the weight s_{DM}. The Decision Maker knows she will able to delude herself into thinking that she observed signal 1, and that she will therefore thinks that project 1 is the right course of action with probability α. Because she can't observe the Implementer's beliefs, the Decision Maker will naturally expect him to have selected his ex post beliefs after having observed signal 2, and therefore expect that he puts in high effort with probability $F(\alpha b_1)$.[6]

Proposition 2 shows that reactivity cannot take place unless both the Decision Maker and Implementer hold realistic beliefs. The above paragraph implicitly assumed that (1) the Decision Maker always chooses to believe in the order she is giving and (2) the Implementer is chooses to be realistic in the reactive equilibrium. We show that these two intuitive

[5]When the signal is 1, the DM always orders 1, by virtue of assumption (1).

[6]After observing signal 2, in a reactive equilibrium, it is optimal for the Implementer to remember signal 2.

strategies are true in equilibrium in Appendix 7.1. The intuitions are the following. First, it is always efficient for the Decision Maker to believe in the order she plans on giving. There are anticipatory gains of doing so (giving the order which has the highest subjective probability of success), but no cost since the decision about the order is chosen simultaneously (so that there is no distortion involved). Second, realism is the optimal decision for the Implementer in an equilibrium where the Decision Maker reacts to the signal. For instance, when the true signal is 2, project 2 will be selected. If the Implementer chooses to believe that in signal 1, she will lower her anticipatory utility, but also provide too little effort since she will think that the selected project is likely to fail.

Finally, one interesting byproduct of proposition 2 is that the reactivity condition is independent of the Implementer's wishful thinking parameter s_I. When the DM gives reactive orders, it is never efficient for the Implementer to be delusional, even when he can enjoy a strong utility from "savoring". This is part of the complementarity highlighted in the introduction: realistic Decision Makers force Implementer to be realistic too. The effect of Implementers on Decision Maker will manifest itself in non reactive equilibria, to which we now turn.

3.3 Non Reactive Equilibria

Non reactive equilibria are equilibria where the Decision Maker always select project 1.

Proposition 3 *Partial Implementer delusion in non reactive organizations:*

Let $\alpha^* \in]1-\alpha;\alpha]$ be the solution of the following optimization problem

$$\max_{x\in[1-\alpha;\alpha]} \int_0^{xb_1} [(s_I x + (1-\alpha))b_1 - (s_I+1)c] dF(c)$$

Then, a non reactive equilibrium exists if and only if:

$$(\alpha s_{DM} F(\alpha b_1) + (1-\alpha) F(\alpha^* b_1))\, B_1 \geq (s_{DM}+1)\, \alpha F((1-\alpha^*) b_2) B_2 \tag{3}$$

In particular, in such an equilibrium:

- *For all signals σ, the Decision Maker remembers signal* 1 *with probability* 1*.*
- *When the true signal is 2, the Implementer remembers signal* 1 *with probability* $\frac{(\alpha^*+\alpha-1)}{2\alpha-1}$*.*

  When the true signal is 1, the Implementer remembers signal 1 *with probability* 1*.*

Proof. See Appendix 7.2. ■

Quite intuitively, the Implementer always remembers project 1 after observing signal 1, as he anticipates the Decision Maker will select project 1. When the true signal is 2, α^* represents the ex post belief held by the Implementer on project 1. If $\alpha^* = \alpha$, he will be completely ignoring the signal and will think 1 is the right project with probability α. If $\alpha^* = 1-\alpha$, the Implementer will remember signal 2 perfectly. The equation defining α^* balances two countervailing forces. On the one hand, increasing α^* increases anticipatory utility by increasing the perceived probability of success. On the other hand, an increase in α^* leads the Implementer to exert inefficient effort. When facing this Implementer's strategy, the Decision Maker can either select project 1 or 2 after observing signal 2. If she selects project 1 (and optimally believes in signal 1), then her realistic self knows that the Implementer's probability of implementing project 1 will be $F(\alpha^* b_1)$ and that project 1 probability of success is $1-\alpha$. Yet, her future self will believe she received signal 1 and thus will believe the Implementer optimally chose to remember signal 1, so that the future (delusional) self should expect a probability of success $\alpha F(\alpha b_1)$. If the Decision Maker deviates and selects project 2 (and optimally chooses to remember signal 2), then both her realistic and future self will expect the Implementer to believe in project 2 with probability $1-\alpha^*$ and will therefore anticipate a probability of success $\alpha F((1-\alpha^*)b_2)$. Condition 3 then simply imposes that the Decision Maker prefers selecting project 1 after observing signal 2. It is immediate to see that as long as condition 3 is verified, it is always optimal for the Decision Maker to select project 1 after observing signal 1.

This proposition leads to several interesting observations. First, the absence of reactivity implies that the Decision Maker becomes fully delusional, while the Implementer is at least partly so: $\alpha^* > 1-\alpha$. The intuition for the Decision Maker is similar to that in the reactive equilibrium: she has nothing to loose in believing in the project she selects. The intuition for the Implementer is slightly different: if the Implementer was completely realistic, he could make a first order gain (on the savoring part of his utility) by increasing his belief in project 1, while making a second order loss only on the realistic part of his utility. This is because this loss would come from the a distorted choice of effort from the viewpoint of her realistic self, but since this effort was optimally set, the effect of a variation of this effort on his utility would be of second order.

It is also interesting to remark that $\alpha^* > 1-\alpha$: the fact that the Implementer will delude himself at least partly facilitates non reactivity by the Decision Maker. This is apparent from condition 3, which is easier to satisfy the larger α^*. If the Implementer believes "more" in project 1, selecting project 2 becomes more costly for the Decision Maker. This is the other side of the complementarity highlighted in the introduction: Implementer's delusion leads the Decision Maker to select the project the Implementer believes is the right one, and therefore makes the Decision Maker delusional herself.

Let us end this section with a last comment. We have focused so far on equilibria that are either fully reactive or fully non reactive. This does not span all potential equilibria. In particular, there are equilibria where projects are selected using mixed strategies as in Landier et al (2009). This is made only to shorten the exposition but does not conceal any important intuition.

4 Comparative Statics

4.1 Wishful Thinking and Reactivity

Proposition 4 *The Scope of Reactivity and the extent of Wishful Thinking*

1. As s_{DM} increases, the scope for non reactivity increases.

2. As s_I increases, (1) the scope for reactivity does not *decrease and (2) the scope for non reactivity increases.*

Proof. See Appendix 7.3. ■

This proposition follows from the conditions described in propositions 2 and 3. The first part of the proposition is straightforward. The second part highlights the complementarity between the Decision Maker and the Implementer's beliefs (a similar point, albeit in a different organizational setting, is made in Bénabou, 2008). When the Decision Maker is realistic and reactive, the Implementer will never want to ignore the truth, whatever her bias. However, when the Implementer has a stronger propensity to delude himself, it tends to make the Decision Maker more non-reactive.

In Landier et al. (2009), we argued that dissenting preferences along the chain of command acted as an efficient disciplining device for the Decision Maker, as it constrained her to use more of the objective, outside information. An interpretation of the above proposition is that wishful thinking represents a limit to the benefits of dissent in organizations, and *that wishful thinking by the Decision Maker has the strongest negative effect when it comes to implementing reactivity.* Dissent in this context can be thought as increasing $b_2 - b_1$ while keeping $b_2 + b_1$ constant. Assume that the organization's Owner has selected a level of dissent such that condition 2 is verified and the organization is reactive for given s_{DM} and s_I. Consider an increase in s_{DM}. If this increase is large enough, it may happen that (1) the organization ceases to be reactive, i.e. condition 2 is no longer verified, and (2) it becomes non reactive, i.e. condition 3 becomes satisfied. If, alternatively, s_I increases, the organization will always remain reactive.

4.2 Multiplicity

Proposition 5 *Multiplicity of Equilibria and Implementer Realism*

Assume that:

$$s_{DM} < \frac{F(b_2) B_2}{F(b_1) B_1 - F(b_2) B_2} \tag{4}$$

1. Assume $s_I > 1$. Then, there exists $\overline{\alpha}$ such that, for all $\alpha > \overline{\alpha}$, both reactive and non reactive equilibria coexist.

2. There exists $\underline{s}$ such that, for all $s_I < \underline{s}$, reactive and non reactive equilibria cannot coexist for any $\alpha \in [1/2, 1]$.

The inequality in this proposition ensures that, for $\alpha = 1$, there exists a reactive equilibrium. This may not happen because when $\alpha = 1$, the Decision Maker receives a very high anticipatory utility from being non reactive: she believes that the project is very likely to succeed, but also believes that the Implementer will agree.

Under condition 4, the above proposition highlights the key complementarity at work in our model. When the propensity to manipulate one's belief is small, we are back to the fully rational case explored in Landier et al (2009). Reactive and non reactive equilibria cannot coexist for a given set of parameters. When, however, the propensity to ignore reality is sufficiently high, the strategic complementarity between the Decision Maker and the Implementer generates multiple equilibria, in particular when α is in the neighborhood of 1. A Decision Maker facing a delusional Implementer will be tempted to always select project 1, while non-reactivity by the Decision Maker will make the Implementer more prone to delusion. The outcome of this two way feedback is that non reactive organizations, provided $s_I > 1$, are sustainable for all levels of reactivity α. Hence, reactivity and non reactivity coexist wherever reactivity exists.

Such multiple equilibria may provide a rationale while a well functioning risk management process suddenly degenerates into collective delusion and becomes unable to interpret warning signs coming from the outside. For instance, it may explain why seasoned buy side investors seemed to rely too much on ratings in their assessment of risk (UBS Report on Write Downs, p 39): both traders and risk managers managed to convince themselves that fixed income securities were not risky anymore. Existing academics analyses focus on ratings shopping (Skreta and Veldkamp, 2008) or suboptimal demand for information (Farhi, Lerner and Tirole, 2008): such models have in common that they can explain why ratings can be overinflated or inaccurate, but do not explain why buy side investors would believe in them. Our model provides one potential explanation..

5 Concluding Remarks

This paper looks at a hierarchy where the Decision Maker and an Implementer have diverging preferences. In normal times, such divergence has the effect of generating reactivity in the organization, but if both agents can select their beliefs, the organizationmay fall in a stable, non reactive equilibrium where all outside information is ignored (very much as in Bénabou, 2008). In the present paper, we have totally ignored issues related to communication and asymmetric information, since all agents observe the signal. If the signal were privately observed by the Decision Maker, then the Implementer would try to interprete the order received. But such interpretation would be impaired by the fact that (1) the DM may have chosen non realistic beliefs and (2) that the Implementer himself may prefer to ignore the truth. This would probably make communication harder, and reactivity less sustainable, in the spirit of the cheap talk model of Dessein (2002). Exploring such interactions further is an interesting lead for future research.

6 References

Bénabou, Roland and Tirole, Jean, 2002, "Self Confidence and Personal Motivation", *The Quarterly Journal of Economics*, 117(3), pp 871-915

Bénabou, Roland, 2008, "Groupthink", *mimeo Princeton University*

Brunnermeier, Markus and Parker, Jonathan, 2005, "Optimal Expectation", *The American Economic Review*, 95(4), pp 1092-1118

Dash, Eric, and Julie Creswell, 2008, "Citigroup saw no Red Flags Even as it Made Bolder Bets", *The New York Times*

Dessein, Wouter, 2002, "Authority and Communication in Organizations", *The Review of Economic Studies*, vol 69, pp 811-838

Farhi, Emmanuel, Lerner, Josh and Tirole, Jean, 2008, "Fear of Rejection? Tiered Certification and Transparency", *mimeo Toulouse School of Economics*

Landier, Augustin, 2000, "Wishful Thinking: A Model of Optimal Reality Denial", mimeo MIT

Landier, Augustin, Sraer, David and Thesmar, David, 2009, "Optimal Dissent in Organizations", *The Review of Economic Studies*

PriceWaterhouseCoopers, 2007, "Creating Value: Effective Risk Management in Financial Services", available at http://www.pwc.com/

Saunders, Anthony, and Cornett, Marcia M., 2005, "Financial Institutions Management: A Risk Management Approach", McGraw-Hill

Skreta and Veldkamp, 2008, "Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation", *prepared for the Carnegie-Rochester Conference Series*

Stulz, René, 2008, "Risk Management Failures: What Are They and When Do They Happen?", *Fisher College of Business Working Papers Series*

UBS Shareholder Report on Write Downs, April 2008

Van Den Steen, Eric, 2007, "The Cost of Incentives Under Disagreement", *mimeo Harvard Business School*

7 Proofs

7.1 Proof of Proposition 2

Consider first the Implementer's incentive problem. Assume the true signal is 2. At date 2, the Implementer expects the Decision Maker to play her equilibrium strategy, i.e. to select project 2. Call Therefore, the Implementer simply maximizes:

$$\max_{x\in[1-\alpha,\alpha]} \int_0^{xb_2} \{(s_I x + \alpha)\, b_2 - (s_I + 1)c\}\, f(c)dc,$$

which admits a corner solution $x = \alpha$.

Similarly, when the true signal is 1, the Implementer expects the Decision Maker to select project 1 and therefore maximizes:

$$\max_{x\in[1-\alpha,\alpha]} \int_0^{xb_1} \{(s_I x + \alpha)\, b_1 - (s_I + 1)c\}\, f(c)dc,$$

which also admits a corner solution $x = \alpha$. Therefore, conditional on the Decision Maker being reactive to signals, it is optimal for the Implementer to fully remember the signals.

Consider now the Decision Maker's incentives. We first remark that the Decision Maker's recollection strategy is simple: it is always optimal for her to think that she received the signal indicating the project she selects at equilibrium. Indeed, assume that there is an equilibrium in which the Decision Maker selects project i, but believes project 1 to be the right project with probability $x < \alpha$. Call β the true probability that 1 is the right project. The Decision Maker's utility then writes: $(s_{DM} x + \beta)\, F(\rho_I b_i) B_i$ and it is clearly optimal for the DM to select $x = \alpha$, i.e. fully believe in project i. This result arises because once a project is selected, there is no cost for the Decision Maker to fully believe in this project.

Therefore, we conclude from the previous analysis that in a reactive equilibrium, the Decision Maker necessarily holds realistic beliefs. The condition for reacting to signal 2 simply states that conditional on signal 2 being the true signal, selecting project 2 (and fully believing in signal 2) delivers a higher utility than selecting project 1 (and optimally fully believing in project 1):

$$(s_{DM}\alpha + \alpha)\, F(\alpha b_2)B_2 \geq (s_{DM}\alpha F(\alpha b_1) + (1-\alpha)F((1-\alpha)b_1))\, B_1$$

Note that the in the previous condition, the Decision Maker holds the Implementer's (realistic) belief as constant, as this is a Nash equilibrium. Also remark that when deviating, the savoring part of the utility will believe in signal 1, and will also believe that the Implementer believes in signal 1. Finally, we simply remark that if the Decision Maker expects the Implementer to be realistic, then it is always optimal to react to signal 1 as using condition 1:

$$\begin{aligned}
(s_{DM}\alpha F(\alpha b_2) + (1-\alpha)F((1-\alpha)b_2))\, B_2 &\leq (s_{DM}+1)\,\alpha F(\alpha b_2)B_2 && (5)\\
&\leq (s_{DM}+1)\,\alpha F(b_2)B_2 && (6)\\
&\leq (s_{DM}+1)\,\alpha F(\frac{b_1}{2})B_1 && (7)\\
&\leq (s_{DM}+1)\,F(\alpha b_1)B_1, && (8)
\end{aligned}$$

which proves that the Decision Maker is better off selecting and remembering project 1 rather than forging a memory for signal 2 and selecting project 2.

7.2 Proof of Proposition 3

Consider a non-reactive equilibrium, i.e. an equilibrium where the Implementer always select project 1. Assume the true signal is 2. The Implementer expects the Decision Maker to select project 1 and will thus maximize:

$$\max_{x\in[1-\alpha,\alpha]} \int_0^{xb_1} \{(s_I x + 1 - \alpha)\, b_1 - (s_I + 1)c\}\, f(c)dc$$

Call $\alpha^\star$ the x maximizing the Implementer's expected utility. It is direct to see that $\alpha^\star > 1-\alpha$ as the derivative in $1-\alpha$ is strictly positive. $\alpha^\star$ is thus either a corner solution (i.e. $\alpha^\star = \alpha$) or an interior solution, in which case, it is defined by the following first order condition:

$$\frac{f(xb_1)}{F(xb_1)} = \frac{s_I}{b_1(x-(1-\alpha))}$$

Note that to hold belief $\alpha^\star$ on project 1 after observing signal 1, the Implementer must forget project 1 with probability ρ such that $\rho\alpha + (1-\rho)(1-\alpha) = \alpha^\star$.

Assume now that the true signal is 1. Then, the Implementer ex post belief on project 1 maximizes:

$$\max_{x\in[1-\alpha,\alpha]} \int_0^{xb_1} \{(s_I x + \alpha)\, b_1 - (s_I + 1)c\}\, f(c)dc$$

which is optimum for $x = \alpha$, i.e. the Implementer fully remembers signal 1 when the true signal is 1.

We now turn to the Decision Maker's incentives. Assume the true signal is 2. First, as we showed in Appendix 7.1, the Decision Maker always choose to remember the signal corresponding to the project she selects. In a non-reactive equilibrium, this amounts to always remembering project 1. Moreover, the Decision Maker should expect ex ante that the Implementer will hold belief $\alpha^\star$ on the probability that 1 is the right project. However, ex post, at equilibrium, the Decision Maker will always remember 1 and will therefore think that the Implementer selected his belief after observing signal 1, and thus optimally chose to remember signal 1. Therefore, the equilibrium utility for the Decision Maker in a non-reactive equilibrium is given by:

$$(s_{DM}\alpha F(\alpha b_1) + (1-\alpha)F(\alpha^\star b_1))\, B_1$$

Assume now that the Decision Maker were to deviate and select project 2 after observing signal 2. Then, ex post, she would optimally remember signal 2. Therefore, she would believe that the Implementer chose his ex post belief after observing 2 and that he therefore holds belief $\alpha^\star$ on project 1. Ex ante, the Decision Maker also believes that the Implementer will hold belief $\alpha^\star$ on project 1. Therefore, the expected utility from deviation for the Decision Maker is given by:

$$(s_{DM}+1)\,\alpha F((1-\alpha^\star)b_2)B_2$$

Therefore, non-reactivity to signal 2 implies that:

$$(s_{DM}\alpha F(\alpha b_1) + (1-\alpha)F(\alpha^\star b_1))\, B_1 \geq (s_{DM}+1)\,\alpha F((1-\alpha^\star)b_2)B_2 \tag{9}$$

Assume now that the true signal is 1. Then, the Decision Maker selects project 1 if:

$$(s_{DM}\alpha+1))\, F(\alpha b_1)B_1 \geq (s_{DM}\alpha F((1-\alpha^\star)b_2) + (1-\alpha)F((1-\alpha)b_2))\, B_2$$

And as soon as condition 9 is verified, the previous condition is verified, so that the Decision Maker always select project 1 after observing signal 1.

7.3 Proof of Proposition 4

We start with the first point of proposition 4.

The derivative of condition 3 with respect to s_{DM} is $\alpha(F(\alpha b_1)B_1 - F((1-\alpha^\star)b_2)B_2)$. Thus, condition 3 is increasing with s_{DM} thanks to condition 1.

The second point of proposition 4 is straightforward. Condition 3 does not depend on s_I. However, it is straightforward to see that α^* is a weakly increasing function of s_I. Indeed, either $\alpha^\star = \alpha$ or $\alpha^\star$ is given by the first order derivative of the maximand in proposition 3, i.e.:

$$G'(x) = 0 = b_1. \{s_I.F(xb_1) + b_1.(1-\alpha-x) f(xb_1)\}$$

which is obviously increasing in s_I. Finally, condition 3 is clearly increasing with α^* and therefore increasing with s_I. **QED**

7.4 Proof of Proposition 5

We start with the first point of Proposition 5. We first show that $\alpha^\star = \alpha$ if $s_I > 1$. Remember that the first order derivative of the Implementer's utility in the no-reactivity equilibrium is given by:

$$\begin{aligned} G'(x) &= b_1. \{s_I.F(xb_1) + b_1.(1-\alpha-x) F'(xb_1)\} \\ G''(x) &= b_1^2. \{(s_I - 1).F'(xb_1) + b_1.(1-\alpha-x) F''(xb_1)\} \end{aligned}$$

given that $F'' < 0$, $F' > 0$, and $x > 1-\alpha$, we deduct that $G'' > 0$ if $s_I > 1$. Therefore, when $s_I > 1$ and since $G'(1-\alpha) > 0$, we conclude that G is monotone (increasing) convex and thus admits a maximum in $\alpha = \alpha^*$.

In this case, the non reactive equilibrium exists if:

$$(1+s_{DM})\,\alpha F((1-\alpha)b_2)B_2 \leq ((1-\alpha)F(\alpha b_1) + s_{DM}\alpha F(\alpha b_1))\,B_1$$

which holds if α close to 1. Reactive equilibria also exist in the neighborhood of $\alpha = 1$ thanks to assumption 4. This proves the first point in Proposition 4. **QED**

The second point of proposition 5 is seen by assuming $s_I = 0$. In this case, $\alpha^* = 1-\alpha$ – it is optimal for the Implementer to fully remember the signal and the conditions of reactivity and non-reactivity becomes, respectively:

$$
\begin{aligned}
(1+s_{DM})\,\alpha F(\alpha b_2)B_2 &\geq ((1-\alpha)F((1-\alpha)b_1)+s_{DM}\alpha F(\alpha b_1))\,B_1 \\
(1+s_{DM})\,\alpha F(\alpha b_2)B_2 &\leq ((1-\alpha)F((1-\alpha)\,b_1)+s_{DM}\alpha F((1-\alpha)\,b_1))\,B_1
\end{aligned}
$$

given that $\alpha > 1/2$, these two conditions are mutually exclusive. The second point of the proposition obtains by continuity. **QED**