How do bonus cap and clawback affect risk and effort choice?: Insight from a lab experiment

By Qun Harris[1], Analise Mercieca[1], Emma Soane[2] and Misa Tanaka[1]

A preliminary draft as of 19 October 2017.

Abstract

We conduct a lab experiment to examine how bonus caps and malus affect individuals' choices of risk and effort. Consistent with the received wisdom, we find that proportional bonus encourages risk-taking, while bonus cap and malus mitigate risk-taking. However, the difference in risk-taking between the bonus cap and malus treatment groups and the proportional bonus group weakened significantly when the participants' bonus was made conditional on hitting an absolute or relative performance target. We also find some evidence that the bonus cap discourages project search effort relative to the proportional bonus, whereas the difference in the levels of effort between the malus group and the proportional bonus group was not statistically significant.

1. Introduction

One of the ironies of the 2007-8 global financial crisis was that senior employees of those banks that have been bailed out by taxpayers walked out of it with their wealth – accumulated through generous bonuses paid up to that point – largely intact. Financial regulators across the world have since reached a consensus that 'compensation practices at large financial institutions are one factor among many that contributed to the financial crisis' (Financial Stability Forum, 2009). In response, a number of jurisdictions have introduced compensation regulations, with the aim of discouraging excessive risk-taking and short-termism and encouraging more effective risk management. In the European Union (EU), the so-called 'bonus cap' was introduced for the 'material risk-takers' at banks, restricting their variable pay to be no more than 100% of the fixed pay (or 200% with shareholder approval).[3] A proportion of the bonus also needs to be deferred, and is made subject to 'malus', implying that it could be forfeited if certain conditions materialise before the deferred pay vests. In addition, a clawback rule was introduced in the United Kingdom, requiring that at least 40% of affected bankers' variable pay is deferred for a period of three to seven years, and that their variable pay can be clawed back for a period of seven to ten years.[4]

Given that these regulations on pay are new and applied exclusively to the employees of banking institutions, it is important to assess whether they achieve their intended aims of mitigating excessive risk-taking, *without* causing significant unintended, detrimental consequences. The theoretical literature predicts that pay structure can affect that both risk and effort choice (e.g.

---

[1] Bank of England, qun.harris@bankofengland.co.uk; analise.mercieca@bankofengland.co.uk; misa.tanaka@bankofengland.co.uk (corresponding author). The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England, the Monetary Policy Committee, the Financial Policy Committee or the Prudential Regulation Committee.

[2] London School of Economics, E.C.Soane@lse.ac.uk.

[3] For the EU bonus cap rules, see DIRECTIVE 2013/36/EU.

[4] For the UK, see the Policy Statement PRA12/15 FCA PS15/16; Remuneration Part of the PRA Rule Book; and the Supervisory Statement on Remuneration, SS 2/17. The rules apply to Material Risk Takers at proportionality level 1 and 2 firms.

Acknowledgement: We are grateful to Lisa Auffegger for research assistance for excellent research assistance. We would also like to thank John Thanassoulis and the attendees of the Bank of England seminar for helpful comments on an earlier version.

Edmans and Liu (2011); Hakanes and Schnabel (2014)), suggesting that pay regulations aimed at curbing excessive risk-taking could also have an unintended impact on effort. The experimental literature also suggests that incentive pay often induces responses other than those for which it was intended, suggesting that changing the structure of pay through regulations could be subject to side effects (see Levitt and Neckermann 2014 for a review).

Yet, empirical identification of the impact of pay regulations on bankers' choices and behaviour is extremely difficult for a number of reasons. First, the data that link individuals' pay with their choices or performances are unavailable, even to the regulators. Second, an empirical study linking firm performance with employee pay contracts may not necessarily enable identification of the impact of pay regulations, as actual pay contracts are likely to be endogenous to unobservable firm, industry and executive characteristics, which in turn could influence observed firm performance (Edmans, Gabaix and Jenter, 2017). Finally, a field experiment with the so-called 'material risk takers' would be impractical, as they normally earn more than €500,000 per annum and hence it is financially too expensive to incentivise them to participate or to influence their behaviour.

We therefore conduct a lab experiment in order to examine how the imposition of specific constraints on bonus – such as bonus cap and malus – could affect risk and effort choices. The experiment was conducted in a behavioural research lab in two separate stages, with a total of 392 participants. The participants engaged in a set of four paid investment tasks, and were paid cash bonus according to the random realisation of the return on the investment asset that they have selected. The experiment, which is structured as a randomised controlled trial, allows us to examine participants' behaviour under three distinct bonus schemes: 1) proportional bonus (representing the 'unregulated' benchmark), 2) capped bonus, and 3) malus. This design allows us to examine how participants' risk and effort choices are influenced by specific 'regulatory' constraints on bonus, and thus enables us to extract qualitative evidence on how remuneration regulations might influence bankers' incentives in real life.

The existing theoretical literature offers a number of hypotheses on the impact of pay regulations on risk and effort choices. Hakanes and Schnabel (2014) have argued that a bonus cap can reduce bankers' risk-taking, but it can also result in suboptimal level of effort. They conclude that the bonus cap is optimal only when the probability of bank bailout is high and the bankers have strong incentives to take excessive risks at the expense of taxpayers in the absence of a bonus cap. Thanassoulis and Tanaka (2017) also argue that malus and clawback regulations can, *ceteris paribus*, mitigate bank executives' excessive risk-taking which is created by the possibility of a bank bailout. But they also argue that imposing malus and clawback through bonus regulations may not, by itself, offer a solution, because the bank owners could respond to these regulations by altering pay contracts so as to undermine their risk-mitigating effects. More specifically, offering the bank executives pay that is convex in the bank's market value can easily restore their risk-taking incentives.

While the lack of data makes it difficult to empirically identify the impact of pay regulations on bankers' incentives and behaviour, there are few studies based on field and lab experiments that give clues on how these regulations might work. For example, Cole, Kanz and Klapper (2015) conduct a framed field experiment with commercial bank loan officers in India, in order to test how different compensation structures affect lending decisions. They find that a bonus scheme that both rewards profitable lending decisions and penalises unprofitable lending decisions (which they call ‘high-powered incentives’) lead to greater effort to screen loan applications than a bonus scheme which simply rewards profitable lending decisions (which they call ‘low-powered incentives’). Effort in their study was measured as the amount of money that loan officers spent on acquiring more detailed information about the loan application. At the same time, they also find that pay deferral – implemented by a 90-day delay in payment – significantly reduced the impact of bonus pay, with dramatically lower effort and profitability of lending.

In order to test the existing hypotheses about the impact of pay regulations on incentives and behaviour, this paper examines how restrictions on bonus that are designed to mimic bonus caps and malus affect individuals’ risk and effort choice in a lab experiment. In our lab experiment, the participants were randomly allocated into three groups that differ in the bonus formula according to which they are paid: i) the control group, representing the ‘unregulated’ benchmark, ii) the bonus cap group, and iii) the malus group. They were then given a series of tasks that require them to choose an investment asset from a list of options that vary in terms of risk-return characteristics. The experiment was designed to examine the following specific hypotheses. First, does a bonus that is proportional on realised investment returns encourage individuals to take greater risks than they would take if they had to invest their own money? Second, do bonus cap and malus mitigate risk taking, relative to proportional bonus? Third, does the risk-taking effect of bonus cap and malus change once bonus is made conditional on meeting an absolute or relative performance target? Finally, do bonus regulations reduce incentives to engage in project search effort?

Our findings suggest that, consistent with the received wisdom, proportional bonus encourages risk-taking, and that bonus cap and malus both mitigate this risk taking. However, this risk-mitigating effect of bonus cap and malus disappear once the bonus is made more convex, e.g. by making the bonus pay conditional on hitting an absolute earnings target, or outperforming a competitor. Thus, our findings are consistent with the hypothesis advanced by Thanassoulis and Tanaka (2017), that the risk-mitigating effects of bonus regulations are easily undermined by bank shareholders through conventionally used incentive tools, such as earnings target and relative performance pay. We also find some evidence that, consistent with Hakenes and Schnabel’s (2014) hypothesis, bonus cap reduces effort. However, we found that malus did not have a statistically significant impact on effort.

Our research contributes to the existing literature in a number of ways. First, our study systematically examines how restrictions on bonus pay influence both risk-taking and project search effort. Our results are largely consistent with the hypothesis of Hakenes and Schnabel (2014): there was strong evidence that bonus cap reduces risk-taking, and some evidence that it also reduces effort. In addition, we also found evidence that malus, which reduces the probability of receiving bonus pay when choosing risky projects relative to the proportional bonus, reduces risk-taking. But we found no evidence that malus reduces effort.

Second, we examine specifically whether it is possible to dilute the impact of the pay regulations by tweaking parameters that are under banks’ control, such as absolute and relative performance metrics. As these are commonly used by banks in rewarding their top executives, it is important to understand whether they can be used as tools to undermine the risk-mitigating effect of pay regulations. For example, amongst large UK banks, the incentive pay of HSBC, Barclays and Standard Chartered CEOs had explicit return-on-equity (RoE) thresholds (e.g. for the vesting of long-term incentive pay) and explicit links to performance relative to a specified peer group. Although there are other experimental studies that have examined how specific restrictions on bonus pay affect risk taking (e.g. Hartmann and Slapnicar (2014)), or how high-powered incentives influence risk shifting to others (e.g. Andersson, Holm, Tyran, and Wengström (2013)), we are not aware of any experimental study which has examined the specific question of how robust the pay regulations are. Our findings are consistent with the hypothesis advanced by Thanassoulis and Tanaka (2017) that the risk-mitigating effect of bonus regulations – both bonus cap and malus – is easily undermined if banks can make executives' pay more convex in response.

The rest of the paper is organised as follows. Section 2 explains the design of our lab experiment. Section 3 presents our findings on how bonus cap and malus affect risk choice. Section 4 presents our results on how these pay restrictions affect effort choice. Section 5 discusses the policy implications and Section 6 concludes.

2. The experimental design

A. The set up

The experiment was conducted in a behavioural research lab in two stages: Study 1 was conducted in January 2017, and Study 2 was conducted in March 2017.

The participants were recruited via the offices of MSc Finance and MSc Management programmes, in order to capture the likely characteristics of 'banker' population as closely as possible in our sample. Thus, most of the participants were students (see Table 1). Participants were recruited with a promise of £10 fee (equivalent to US$12.5 or €11.7 as at end-March 2017) and an opportunity to earn an additional bonus for participating in an hour-long lab experiment. A Math A-levels qualification or equivalent was set as a minimum requirement for participating in the experiment, in order to increase the likelihood that the participants make reasonable choices in investment tasks which require an understanding of probabilities.

After eliminating participants who have followed the instructions incorrectly, we were left with a sample of 219 in Study 1 and 173 in Study 2, so a total of 392. Table 1 reports the sample characteristics.

Table 1: Sample characteristics

		Study 1				Study 2				Whole Sample			
		Total	Control	Bonus Cap	Malus	Total	Control	Bonus Cap	Malus	Total	Control	Bonus Cap	Malus
Number o fparticipants		219	74	74	71	173	57	57	59	392	131	131	130
of which:		100.0%	33.8%	33.8%	32.4%	100.0%	32.9%	32.9%	34.1%	100.0%	33.4%	33.4%	33.2%
Students		92%	92%	91%	93%	79%	84%	72%	81%	86%	89%	82%	88%
Sex	Female	61.2%	56.8%	58.1%	69.0%	61.3%	59.6%	68.4%	55.9%	61.2%	58.0%	62.6%	63.1%
	Male	38.4%	43.2%	40.5%	31.0%	38.2%	40.4%	31.6%	42.4%	38.3%	42.0%	36.6%	36.2%
	Undisclosed	0.5%	0.0%	1.4%	0.0%	0.6%	0.0%	0.0%	1.7%	0.5%	0.0%	0.8%	0.8%
Age	Mean	23.6	23.3	24.4	23.0	26.2	25.0	29.0	24.7	25.3	24.5	27.0	24.5
	Median	22.0	21.0	22.0	22.0	23.0	23.0	24.0	22.0	23.0	22.0	23.0	22.0

Both studies consisted of the following three parts:

Part 1: Questions relating to their personality and risk preference, and a set of probability questions presented as a 'tutorial' (see Annex).

Part 2: Participants were given four tasks in which they were asked to select an asset from a list of options, and were paid according to a specific bonus scheme which was explained to them before engaging in the task. Participants were randomly assigned to three different bonus regimes: i) control group, ii) bonus cap group, and iii) malus group.

Part 3: Questions on biographical information (age, sex, education, work experience, etc).

The participants in both Studies 1 and 2 completed a common set of questions in Parts 1 and 3. However, the design of Part 2 differed in Studies 1 and 2, as explained more in detail below. Participants were paid a bonus only based on the outcome of their choices in Part 2.

B. Hypothetical investment task (Task 0, Part 1)

In order to isolate the impact of the bonus regime on risk choice, we need to control for individuals' risk preference. Thus, Part 1 of both Study 1 and 2 included a series of psychological questionnaire which was aimed at extracting individuals' risk preference.

In addition, Part 1 of both in Studies 1 and 2 included a common *hypothetical* asset choice task, which we call Task 0. In Task 0, participants were given a hypothetical scenario, in which they inherited £100,000 from a distant relative whom they have never met before. They were then asked to choose one asset from the list of six assets, as shown in Table 2: the assets were listed in the order of ascending riskiness (with Asset 1 being risk free and Asset 6 being the most risky), and they differed in risk-return characteristics. In Task 0, columns (c), (d), and (e) were presented in the units of £ thousand. No cash bonus was paid based on the outcome of the investment chosen in this task. In subsequent analysis, we use the participants' choice in this hypothetical task as a variable to control for their inherent risk preference.[5]

Table 2: 6-asset choice task

	Probability of failure	Probability of success	Return when failure	Return when succcess	Expected return
	(a)	(b)	(c)	(d)	(e)
Asset 1	0%	100%	0.0	1.0	1.0
Asset 2	10%	90%	-0.5	2.0	1.8
Asset 3	20%	80%	-1.5	3.0	2.1
Asset 4	30%	70%	-2.5	4.5	2.4
Asset 5	40%	60%	-3.5	6.0	2.2
Asset 6	50%	50%	-4.5	8.0	1.8

Note: In both Studies, columns (c), (d) and (e) were given in units of £ thousand in Task 0, whereas they were given in £ million in all other 6-asset Tasks (Study 1 Tasks 1 and 3, and Study 2 Tasks 1, 3 and 3A).

C. Bonus groups in paid tasks (Part 2)

In both Studies 1 and 2, the participants were asked to engage in four paid investment tasks in Part 2. In each task, the participants were given a scenario in which they had to act as an investment manager at a hypothetical 'ABC Bank', and had to choose one asset to invest in from a list of assets that differed in risk-return characteristics. The participants were told that they could earn a bonus depending on the return on the asset they chose in each task. After completing the study, participants received a cash payment that represented 1/1000 of the hypothetical bonus that they

[5] Kühberger, Schulte-Mecklenbeck and Perner (2002) find that hypothetical choices match real choices for both for small as well as large payoffs.

had earned in Part 2 of the study. For example, if a participant earned a hypothetical bonus of £10,000 in the study, she received a cash bonus of £10 in addition to the £10 participation fee. In order to prevent the possibility that information about the outcome of one investment task influences subsequent choices, the participants were informed of the outcome of their investment and the bonus they have earned only after completing all the tasks.

Before engaging in the four investment tasks in Part 2, the participants were randomly allocated into one of the three bonus regime groups:

Group 1: Control group. Participants were paid a bonus proportional to the asset return, if the return exceeded a pre-specified threshold. No bonus was paid if the return was below the threshold.

Group 2: Bonus cap group. Participants were paid a bonus proportional to the asset return, once the return exceeded a pre-specified threshold, but the maximum bonus a participant could earn in each task was capped at £4,000 (i.e. £4 real cash bonus per task paid at the end of the study). No bonus was paid if the return was below the threshold.

Group 3: Malus group. Participants were paid a bonus proportional to the asset return, once the return exceeded a pre-specified threshold in the first 'year', conditional on the project not failing for another 'year'. The probability of success in each 'year' was set to be the same: For example, in the case of the baseline asset choice task presented in Table 2, the probability of success in each year is given in column (e): so for example, the probability of Asset 6 succeeding for two consecutive years was (50%)^2=25%. No bonus was paid if the return was below the threshold.

There was no framing effect in presenting the bonus schemes. As explained in the next section, the threshold for bonus payment varied across tasks for each group. The actual payment was rounded up to the nearest £0.5 (50 pence).

Table 3: The maximum possible bonus and total pay for different bonus groups

	Fixed participation fee (a)	Maximum possible bonus (b)	Maximum possible total pay (a) + (b)
Group 1 (control)	£10	£33	£43
Group 2 (bonus cap)	£10	£16	£26
Group 3 (malus)	£10	£33	£43

Table 3 shows the maximum bonus and total pay that participants could earn in the four tasks in both studies. For all three bonus groups, the amount of bonus that participants could earn was sufficiently large relative to their fixed participation fee in order to incentivise them to make considered choices. Because the bonus was capped at £4 per task for Group 2 (bonus cap), participants in this group could earn only up to a maximum of £16 in total in the four investment tasks. The level of the bonus cap was calibrated to be consistent with the EU-wide regulation that restricts bankers' variable pay to be no more than 100% of the fixed pay, and no more than 200% of the fixed pay with shareholders' approval: thus, the total bonus that the bonus cap group could earn was calibrated to be 160% of the fixed pay (the participation fee) of £10. The maximum bonus that participants in Group 3 (malus) could earn was as large as that of the control group, but as we explain below, the probability that the malus group participants could earn this amount was much lower than that facing the control group.

The Group 3 bonus scheme is designed to mimic certain features of malus, by which an individual receives a bonus if the investment is successful, but could lose the entire bonus if it fails in the subsequent years. We acknowledge that the design of the Group 3 bonus scheme differs from the typical use of malus in the real world in two ways. First, in reality, bankers very rarely lose their entire unvested bonus even if malus is applied. Second, the Group 3 bonus scheme does not capture the time dimension of bonus deferral, as all participants were paid immediately in cash after the experiment. We opted for this design in order to ensure that the explanation of the bonus scheme is sufficiently simple for the participants to understand; and to focus our analysis on the impact of incentivizing the participants to care about the long-term risk, which is the main aim of the existing pay regulation requiring malus.

D. The design of paid investment tasks (Part 2)

In both studies, participants were asked to engage in four paid investment tasks (Tasks 1-4), in which they had to select one asset from a list given to them. The participants stayed in the same bonus group throughout this part of the study, were given full information about their bonus scheme before engaging in each task.

<u>Study 1</u>

The 219 participants in Study 1 were randomly assigned into three different bonus groups and engaged in the following four paid investment tasks.

Task 1: All participants were asked to choose from the list of 6 assets in Table 2, where the bonus threshold for all bonus groups was set at zero: columns (c), (d) and (e) in Table 2 were presented in units of £ million in this task (as well as in Task 3, which is explained below). That means that participants in Group 1 and 2 could earn a bonus in Task 1 as long as the asset return was positive; and those in Group 3 could earn a bonus as long as the asset return was positive for two consecutive 'years'. Our objective in designing this task was to examine whether different bonus groups make systematically different choices.

Task 2: the bonus scheme for each group remained the same as in Task 1. But now, participants were asked to select from a list of 30 assets, presented in Table 4 below: columns (c), (d) and (e) in Table 4 were presented in units of £ million in this task (as well as in Task 4, which is explained below). The first 6 assets in the list were designed to be the same as in Task 1 (see Table 1). The following 24 assets were grouped in a set of 6 assets, and were designed to be either dominant to, or dominated by, the first 6 assets: compared to Assets 1-6, Assets 7-12 have lower probability of success, Assets 13-18 have a lower payoff, Assets 19-24 have a higher probability of success, and Assets 25-30 have a higher probability of success and higher payoff when success. Of these 30 assets, Assets 19, 25-30 were the 'dominant' assets, in the sense that they offered a better return for the same risk, or lower risk for the same return, than other assets in the list. Identifying these assets would have required cognitive effort of having to scan large numbers by scrolling up and down the screen, but any rational participants should have chosen one asset from this set, regardless of his or her risk preference. Thus, we deemed a participant to have exerted effort only if that person selected an asset from this 'dominant' set (Asset 19, 25-30). Our objective in designing this task was to examine whether the bonus regime influences project search effort.

Table 4: 30-asset choice task (Study 1, Task 2)

	Probability of failure	Probability of success	Return when failure	Return when succcess	Expected return
	(a)	(b)	(c)	(d)	(e)
Asset 1	0%	100%	0.0	1.0	1.0
Asset 2	10%	90%	-0.5	2.0	1.8
Asset 3	20%	80%	-1.5	3.0	2.1
Asset 4	30%	70%	-2.5	4.5	2.4
Asset 5	40%	60%	-3.5	6.0	2.2
Asset 6	50%	50%	-4.5	8.0	1.8
Asset 7	5%	95%	0	1.1	1.0
Asset 8	15%	85%	-0.5	1.5	1.2
Asset 9	25%	75%	-1	2.5	1.6
Asset 10	35%	65%	-2	4	1.9
Asset 11	45%	55%	-3	6	2.0
Asset 12	55%	45%	-4	7.5	1.2
Asset 13	0%	100%	0	0.9	0.9
Asset 14	10%	90%	-0.5	1.9	1.7
Asset 15	20%	80%	-1.5	2.9	2.0
Asset 16	30%	70%	-2.5	4.4	2.3
Asset 17	40%	60%	-3.5	5.9	2.1
Asset 18	50%	50%	-4.5	7.9	1.7
Asset 19	0%	100%	0	1.1	1.1
Asset 20	9%	91%	-0.5	1.5	1.3
Asset 21	19%	81%	-1	2.5	1.8
Asset 22	29%	71%	-2	4	2.3
Asset 23	39%	61%	-3	6	2.5
Asset 24	49%	51%	-4	7.5	1.9
Asset 25	0%	100%	0	1.1	1.1
Asset 26	9%	91%	-0.5	2.3	2.0
Asset 27	19%	81%	-1	3.3	2.5
Asset 28	29%	71%	-2	4.8	2.8
Asset 29	39%	61%	-3	6.3	2.7
Asset 30	49%	51%	-4	8.3	2.3

Note: This table was used in Study 1 Tasks 2 and 4, and Study 2 Task 2A. In all these tasks, columns (c) to (e) were given in units of £ million.

Task 3: The participants were again given the 6-asset choice task as in Task 1 (see Table 2), but now the participants were given an absolute performance target of £5 million of asset return and were told that they will not be paid a bonus unless the asset return is £5 million or greater. Thus, Task 3 raised the bonus threshold from £0 in Task 1 to £5 million. Our objective here was to examine whether any differences in risk choice under different bonus regimes observed in Task 1 remained even when an absolute performance target is introduced and bonus is used to reward only exceptional performance.

Task 4: The participants were again given the 30-asset choice task, but under the same bonus regime as in Task 3. The participants were given a table containing the same 30 assets as in Task 2 (shown in Table 4), but the assets were re-ordered so as to eliminate the incentive to choose the same asset as Task 2 without scrutinizing the characteristics of the assets. The objective here was to examine whether any differences in effort choice under different bonus regimes observed in Task 2 remained even when the bonus is made more convex.

Let p_i, $r_{i,s}$ and R_i be the probability of success, the return when success, and the *realized* return when the participant chose Asset *i*, respectively. In Tables 2 and 4, p_i is given in column (b), $r_{i,s}$ is given in column (d), and R_i takes the value in column (c) if the investment fails, and the value in column (d) if it succeeds. Table 5 provides a summary of the bonus regimes for each task in Study 1: for example, in Task 1, if a participant chose Asset 3, she was paid £3 for the task if the investment was successful and zero otherwise. Note that, for malus group, p_i could be interpreted as short-term risk, whereas p_i^2 can be interpreted as long-term risk of Asset *i*. After completing all four tasks, the bonus payment for each task was generated based on the risk-return characteristics of the asset chosen by the participant, and the bonus group that the participant was assigned to.

Table 5: Bonus regimes in Study 1

	Task 1 (6 assets)	**Task 2 (30 assets)**	**Task 3 (6 assets)**	**Task 4 (30 assets)**
Group 1 (control)				
Probability of bonus payment	p_i	p_i	p_i if $r_{i,s} \geq 5$ 0 otherwise	p_i if $r_{i,s} \geq 5$ 0 otherwise
Amount of bonus	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise
Group 2 (bonus cap)				
Probability of bonus payment	p_i	p_i	p_i if $r_{i,s} \geq 5$ 0 otherwise	p_i if $r_{i,s} \geq 5$ 0 otherwise
Amount of bonus	$Min(R_i, 4)$ if $R_i \geq 0$ 0 otherwise	$Min(R_i, 4)$ if $R_{i,s} \geq 0$ 0 otherwise	4 if $R_i \geq 5$ 0 otherwise	4 if $R_i \geq 5$ 0 otherwise
Group 3 (malus)				
Probability of bonus payment	p_i^2	p_i^2	p_i^2 if $r_{i,s} \geq 5$ 0 otherwise	p_i^2 if $r_{i,s} \geq 5$ 0 otherwise
Amount of bonus	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise

Note: p_i, $r_{i,s}$ and R_i are the probability of success, the return when success, and the *realized* return when the participant chose Asset *i, respectively*. In Tables 2 and 4, p_i is given in column (b), $r_{i,s}$ is given in column (d), and R_i takes the value in column (c) if the investment fails and the value in column (d) if it succeeds.

Study 2:

Study 2 builds on Study 1 to examine two issues: i) whether participants in different bonus groups exert different levels of effort when effort task was made harder, and ii) how the relative performance target influences risk taking under different bonus regimes. As in Study 1, participants were randomly assigned to three bonus groups. A total of 173 participants were included in the sample for Study 2.

In Study 2, Tasks 1 and 3 w exactly the same as in Study 1: this allowed us to examine the key issue of how bonus convexity, introduced via an earnings target, affected the impact of the bonus regulation with an expanded sample. But Study 2 differed from Study 1 in two ways. First, Task 2 in Study 1 was replaced by Task 2A in Study 2.

Task 2A: The participants were shown the first 6 assets in Table 2: columns (c), (d) and (e) in Table 2 were presented in units of £ million in this task. But they were told that, in order to reveal each of the additional 24 assets, they had to perform a 5-digit, 3-number summation or subtraction question correctly (e.g. 30582+28951+49501=?; 48206-17829-12938=?).[6] The participants were given a pencil and paper to perform the calculations manually, and were not allowed to use a calculator or similar devices. The participants were given the opportunity to do up to 24 sets of calculations in order to reveal the full set of 30 assets in Table 4, but they were not informed whether the 24 hidden assets were better or worse than the first 6 assets, so the reward from effort was *ex ante* uncertain. The participants could choose not to do any calculation, or stop doing the calculation at any point in time. So for example, if they attempted 10 questions and answered 8 correctly and stopped at that point, they could see additional 8 assets, so they could see Assets 1 to 14 in Table 4 before making their final asset choice. Participants were informed whether they correctly answered an algebra question, before deciding whether to attempt the next question or stop the algebra task and view the assets revealed.

In this task, the participants' effort was measured both by the number of calculations attempted, which was their choice variable. There were two aims in designing Task 2A in Study 2. First, this task was designed to be harder than Task 2 in Study 1, as the participants had to engage in mentally onerous math computations in order to reveal each asset. At the same time, the math problems were set at levels such that a primary school child can solve them with some effort, and hence did not require learning, high-level math education or ability in order to perform them correctly. Second, unlike Task 2 in Study 1, this task introduced uncertainty in the return from effort. This set-up seeks to capture the reality in which bankers searching for investment opportunities typically do not know *ex ante* whether their effort would lead to a discovery of good projects, and hence make their effort choices under uncertainty.

Second, Task 4 in Study 1 was replaced by Task 3A in Study 2.

Task 3A: The participants were asked to choose from the list of six assets shown in Table 2, but were told that, in order to receive a bonus, their asset return has to be as high as, or higher than that of a fictitious competitor. The participants were also told that the competitor has chosen Asset 5 (see Table 2). Columns (c), (d) and (e) in Table 2 were presented in units of £ million in this task.

As in Study 1, the participants in Study 2 did not find out the return on their assets or the bonus earned until they completed all four tasks. The payment was generated using the same method as in Study 1, except in Task 3A where the payment depended both on the randomly generated return on the asset chosen by the participant, and the randomly generated return on Asset 5 which was chosen by the fictitious competitor.

As before, we denote the probability of success, the return when success, and the *realized* return as p_i, $r_{i,s}$ and R_i, respectively, when when the participant chooses Asset *i*. In addition, we denote as ρ_5 the *realized* return on the competitor's chosen Asset 5 in Task 3A, which can take the value of either -3.5 or 6.0 (see Table 2). The mathematical expressions of the bonus regimes in Study 2 are summarized in Table 6. Note that the bonus regimes in Tasks 1, 2A, and 3 in Study 2 were the same as those in Tasks 1, 2, and 3 in Study 1. In Task 3A, the participants were paid only if their investment succeeded and their competitor's investment failed (which occurred with a 40% probability), if they had chosen assets that would yield less than £5 million even when the

[6] Performance of mathematical calculations were used as 'real effort task' to measure the level of effort in previous studies, e.g. Brüggen and Strobel (2007), Dohmen and Falk (2011), Eriksson, Poulsen and Villeval (2009).

investment succeeded (i.e. Assets 1-4). By contrast, they were paid as long as their investment succeeded, *regardless of whether the competitors' investment succeeded or not*, if they had chosen Asset 5 or 6 which would yield a return as high as, or higher than the competitor's investment (Asset 5) if it was to succeed.

Table 6: Bonus regimes in Study 2

	Task 1 (6 assets)	**Task 2A (30 assets)**	**Task 3 (6 assets)**	**Task 3A (6 assets)**
Group 1 (control)				
Probability of bonus payment	p_i	p_i	p_i if $r_{i,s} \geq 5$ 0 otherwise	$0.4*p_i$ if $r_{i,s} < 6$ p_i if $r_{i,s} \geq 6$
Amount of bonus	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise	R_i if $R_i \geq \rho_5$ 0 otherwise
Group 2 (bonus cap)				
Probability of bonus payment	p_i	p_i	p_i if $r_{i,s} \geq 5$ 0 otherwise	$0.4*p_i$ if $r_{i,s} < 6$ p_i if $r_{i,s} \geq 6$
Amount of bonus	$Min(R_i, 4)$ if $R_i \geq 0$ 0 otherwise	$Min(R_i, 4)$ if $R_{i,s} \geq 0$ 0 otherwise	4 if $R_i \geq 5$ 0 otherwise	$Min(R_i, 4)$ if $R_i \geq \rho_5$ 0 otherwise
Group 3 (malus)				
Probability of bonus payment	p_i^2	p_i^2	p_i^2 if $r_{i,s} \geq 5$ 0 otherwise	$0.4*p_i^2$ if $r_{i,s} < 6$ p_i^2 if $r_{i,s} \geq 6$
Amount of bonus	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 0$ 0 otherwise	R_i if $R_i \geq 5$ 0 otherwise	R_i if $R_i \geq \rho_5$ 0 otherwise

Note: p_i, $r_{i,s}$ and R_i are the probability of success, the return when success, and the *realized* return when the participant chose Asset *i, respectively*. In Tables 2 and 4, p_i is given in column (b), $r_{i,s}$ is given in column (d), and R_i takes the value in column (c) if the investment fails and the value in column (d) if it succeeds. ρ_5 is the realised return on the competitor's chosen Asset 5, which could take a value of -3.5 or 6.0 (see columns (c) and (d) in Table 2.)

3. Analysis: The impact of bonus regimes on risk-taking

In order to examine the impact of bonus regimes on risk-taking, we examine participants' asset choices under the common risk choice tasks (Task 0, Task 1 and Task 3) in Studies 1 & 2, and Task3A, which was included only in Study 2 to study the impact of relative performance benchmarking. The participants' choices in these tasks are found in Table A in the Annex.

A. Impact of proportional bonus on risk-taking

In order to first assess the impact of proportional bonus on risk-taking, we compare the choices of the participants assigned to the control group (Group 1) in Studies 1 and 2 in Task 0 and Task 1. A total of 131 participants were assigned to the control group in Study 1 (74 participants) and Study 2 (57 participants). In both Task 0 and Task 1, the participants were asked to choose from the same list of 6 assets in Table 2. In Task 0, the participants were asked to invest a hypothetical inheritance: no cash bonus was paid based on the outcome of this hypothetical task. In Task 1, the participants acted as investment managers for a hypothetical bank and earned a cash bonus depending on (proportional to) the realised return of the asset they chose. We use these two tasks to examine the following hypothesis (H1):

H1: Proportional bonus encourages individuals to take greater risks than they would take with their own money.

For the purpose of statistical analysis, the participants' asset choices were grouped into three risk level categories: Risk level 1 (Assets 1-3), Risk level 2 (Asset 4), and Risk level 3 (Asset 5-6). The rationale for categorising the assets into these three risk bucket is as follows. In the absence of any distortion in incentives, Asset 4 represents a *risk neutral* choice[7], as it offers the maximum expected returns on the bank's investment asset. Assets 1-3 represent *risk averse* choices: they offer lower expected returns on the bank's investment asset relative to Asset 4, but also lower probabilities of failure. Asset 5-6 represent *risk loving* choices: they yield lower expected return on the bank's investment asset than Asset 4, but also higher probabilities of failure.

As shown in Table A in the Annex, 10.8% more participants in the Study 1 control group chose Assets 5-6 (*Risklevel_3*) in Task 1 compared in Task 0; whereas 8.8% more participants in the Study 2 control group chose Assets 5-6 (*Risklevel_3*) in Task 1 compared in Task 0.

To assess the impact of proportional bonus on risk-taking, we used the following maximum-likelihood multinomial logit models with discrete dependent variables to test the statistical significance of the observed differences between Task 0 and Task 1 choices of those in the control group:

$$ln\frac{Pr(Risklevel = 1)}{Pr(Risklevel = 2)} = C_{10} + \ C_{11}Bonus + C_{12}Male + C_{13}Age$$

$$ln\frac{Pr(Risklevel = 3)}{Pr(Risklevel = 2)} = C_{20} + \ C_{21}Bonus + C_{22}Male + C_{23}Age$$

where the dependent variables were the risk levels chosen in Task 0 and Task 1 (*Risklevel*). The right-hand side variables included a dummy *Bonus*=1 if the asset choice was made in Task 1, and Bonus=0 when in Task 0, and a dummy *Male* =1 if the participant a male and 0 otherwise. We also included *Age* in the regression. Table 7 summarises our results. Consistent with our hypothesis (H1), the *Bonus* dummy was positive and significant for *Risklevel_3*, suggesting that, when the participants were paid a proportional bonus, they were more likely to choose a highly risky asset (*Risklevel_3*) than in a hypothetical scenario in which they were asked to invest their inheritance. We interpret this result as supporting the hypothesis that the proportional bonus increases risk taking. Age was statistically significant: older participants were more likely to choose lowest risk assets (*Risklevel_1*) that Asset 4 (*Risklevel_2*), and older participants were marginally more likely to choose highest risk assets (*Risklevel_3*) than Asset 4 (*Risklevel_2*). Gender was also statistically significant: male participants were more likely than female participants to choose highest risk assets (Risklevel_3) than Asset 4 (*Risklevel_2*).

[7] We label Asset 4 as the risk neutral choice as it represents the optimal choice of risk neutral individuals who seek to maximise the expected return. However, it is possible that some risk averse (or risk loving) individuals will also choose this asset, if they consider the risk-return trade-offs of lower (or higher) risk assets to be unattractive relative to Asset 4.

Table 7: Impact of proportional bonus on risk choice

	(1) Risklevel	
Risklevel_1		
Bonus	0.022	(0.271)
Male	-0.000	(0.280)
Age	0.067***	(0.023)
Constant	-1.967***	(0.567)
Risklevel_3		
Bonus	1.196**	(0.478)
Male	0.944**	(0.446)
Age	0.053*	(0.031)
Constant	-4.094***	(0.877)
Observations	262	
Pseudo R^2	0.049	

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B. Impact of bonus cap and malus in mitigating risk-taking

Next, we examine how bonus cap and malus affect the participants' Task 1 risk choices, in order to test the following hypothesis:

H2: Bonus cap and malus mitigate risk taking, relative to proportional bonus.

To test the above hypothesis, we estimated the following maximum-likelihood multinomial logit models with discrete dependent variables (i.e. the participants' choices of assets in *Risklevel*=1, 2 or 3 categories in Task 1), controlling for participants' inherent risk preferences (i.e, their asset choices in Task 0) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(t1\_risklevel = 1)}{Pr(t1\_isklevel = 2)} = C_{10} + C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(t1\_risklevel = 3)}{Pr(t1\_isklevel = 2)} = C_{20} + C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where t1_risklevel denotes the Risk level (=1,2 or 3) the participants chose in Task 1, BonusRegime=2 is a dummy which equals one if the participant is in Group 2 (bonus cap), and BonusRegime=3 is a dummy which equals one if the participant is in Group 3 (malus). Inheritancechoice denotes participants' asset choice in Task 0.

Table 8 summarises our results. We found that, compared with those in the control group, when controlling for Task 0 choice (which captures the participants' inherent risk preference), the participants in the malus group were less likely to choose the assets of the highest risk level (*Risklevel_3*) than Asset 4, and more likely to choose the assets of the lowest risk level (*Risklevel_1*) than Asset 4. Relative to the control group, the participants in the bonus cap group were also less likely to choose the assets of the highest risk level (*Risklevel_3*) than Asset 4. Thus, the evidence is consistent with our hypothesis (H2).

Table 8: The impact of bonus regime on Task 1 risk level choice

	(1) t1_risklevel	
Risklevel_1		
Control	0.000	(.)
Bonus cap	0.335	(0.292)
Malus	0.743^{***}	(0.288)
Inheritancechoice	-0.990^{***}	(0.139)
Constant	2.882^{***}	(0.503)
Risklevel_2		
Control	0.000	(.)
Risklevel_3		
Control	0.000	(.)
Bonus cap	-2.125^{***}	(0.663)
Malus	-2.519^{***}	(0.829)
Inheritancechoice	1.574^{***}	(0.344)
Constant	-7.457^{***}	(1.421)
Observations	392	
Pseudo R^2	0.193	

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

These findings are consistent with what one might expect. Under the malus regime, participants were paid a bonus proportional to the asset return, conditional on the asset's return being positive for another 'year'. Participants became more cautious under this regime: they were both less likely to choose the assets of the highest (=3) risk level and more likely to choose assets of the lowest (=1) risk level. For the bonus cap regime, the participants were paid a bonus proportional to the asset return, subject to a cap. Therefore, relative to the bonus scheme of the control group, the bonus cap reduced the bonuses that the participants could get by investing the higher risk/higher return assets, but not the bonus that they could get by investing in the lower risk/lower return assets (which returns are under the cap). It is therefore not surprising that under the bonuscap regime, the participants were less likely to choose the assets of the highest (=3) risk level, while there was not a statistically significant change in their risk preference among assets of the lowest (=1) risk level.

Inheritancechoice is also found to be statistically significant. The negative coefficient between the relative log probability of choosing *t1_risklevel =1* and *Inheritancechoice* suggests that participants who chose high-risk, high-return assets in Task 0 were less likely to choose assets of the lowest risk level (*Risklevel_1*). The positive coefficient between the relative log probability of choosing t1_risklevel =3 and Inheritance choice suggests that participants who chose high-risk, high-return assets assets in Task 0 were more likely to choose assets of the highest Risklevel (=3).

Next, we examined whether the bonus regime influenced how the risk level chosen in Task 1 has changed from that chosen in Task 0. If the risk level of the asset chosen in Task 1 is higher than that chosen in Task 0, we call it 'risk up'; if it is the same, 'no change'; and if the risk level in Task 1 is lower than in Task 0, then 'risk down'. We estimated the following maximum-likelihood multinomial logit models with discrete dependent variables (i.e., Risk up, No change, Risk down from Task 0 to Task 1) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(i\_task1 = 1)}{Pr(i\_task1 = 2)} = C_{10} + C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(i\_task1 = 3)}{Pr(i\_task1 = 2)} = C_{20} + C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where i_task1 denotes the change of participants' choices in Task 1 and Task 0, with i_task1=1, 2 and 3 corresponding to "risk down", "no change" and "risk up" respectively. BonusRegime =1, 2 or 3 corresponds to control (proportional), bonus cap and malus regimes, and Inheritancechoice denotes the participants' asset choices in Task 0.

Table 9 summarises our results. We found that, compared with those in the control group, the participants in the malus group were statistically less likely to risk up and more likely to risk down, while the participants in the bonus cap group were statistically less likely to risk up. These findings are consistent with the results from the analyses of the risk levels chosen by different bonus groups, supports our hypothesis (H2). The participants in the malus group were more cautious: both less likely to risk up and more likely to risk down. The participants in the bonus cap group were less likely to risk up, as the potential bonus from choosing a higher risk/higher return asset was capped, while there was no definitive evidence of an increase in risk down as the cap did not affect the potential bonus of choosing lower risk/lower return assets.

Table 9: The impact of bonus regime on risking up and down in Task 1, relative to Task 0

	(1) i_task1	
risk_down		
Control	0.000	(.)
Bonus cap	0.525	(0.367)
Malus	1.110***	(0.348)
Inheritancechoice	0.505***	(0.160)
Constant	-3.625***	(0.694)
no_change		
Control	0.000	(.)
risk_up		
Control	0.000	(.)
Bonus cap	-0.747**	(0.351)
Malus	-1.300***	(0.410)
Inheritancechoice	-0.957***	(0.147)
Constant	2.250***	(0.513)
Observations	392	
Pseudo R^2	0.133	

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

These findings are consistent with our hypothesis (H2) that, when compared with the control group (proportional bonus scheme), both bonus cap and malus mitigate greater risk taking, in the absence of any other intervention.

Inheritancechoice was found to be statistically significant. Other things being equal, participants who chose high risk/high return assets in Task 0 (Inheritancechoice) were less likely to risk up. This is because in the experiment setup, the probability of risk up (in Task 1) for a participant who chose Asset 6, the highest risk/return asset among the available choices, in Task 0 was zero. Participants who chose lower risk/lower return assets in Task 0 (Inheritancechoice) were less likely to risk down. This is because in the experiment setup, the probability of risk down (in Task 1) for a participant who chose Asset 1, the lowest risk/return asset among the available choices, in Task 0 would be zero.

C. Impact of an absolute performance target on the risk-mitigating effects of bonus cap and malus

In Task 3 (Studies 1 & 2), participants were again presented the same 6-asset choice task as in Task 1 (see Table 2), but now the participants were given an absolute performance target of £5 million. They would be paid a bonus only when the asset they chose generated a return of £5 million or greater. We now examine the following hypothesis:

H3: The risk-mitigating effect of bonus cap and malus weakens once an absolute earnings target is introduced.

The data in Table A in the Annex show that the introduction of an absolute earning target increased the number of participants who chose the highest Risklevel (=3) assets substantially across all three bonus groups. We fitted data to maximum-likelihood multinomial logit models with discrete dependent variables (i.e., Risklevel=1, 2 or 3), controlling for participants' inherent risk preferences (i.e, their choices in Task 0) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(t3\_risklevel = 1)}{Pr(t3\_isklevel = 3)} = C_{10} + C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(t3\_risklevel = 2)}{Pr(t3\_isklevel = 3)} = C_{20} + C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where t3_risklevel denotes the Risklevel (=1,2 or 3) participants chose in Task 3, BonusRegime (=1, 2 or 3) corresponds to proportional (the control), bonus cap and malus regimes. Inheritancechoice denotes the asset choices participants made in Task 0.

Table 10: Impact of bonus regimes on risk choice, with an earnings target

	(1)	
	t3_risklevel	
Risklevel_1		
Control	0.000	(.)
Bonus cap	-0.964**	(0.454)
Malus	-0.622	(0.429)
Inheritancechoice	-0.383**	(0.168)
Constant	1.945***	(0.641)
Risklevel_2		
Control	0.000	(.)
Risklevel_3		
Control	0.000	(.)
Bonus cap	-0.315	(0.368)
Malus	-0.532	(0.364)
Inheritancechoice	0.118	(0.141)
Constant	1.330**	(0.566)
Observations	392	
Pseudo R^2	0.031	

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We found that, once the earnings target was introduced, bonus cap and malus groups were as likely to choose assets of the highest risk level (*Risk level_3*) as the control group: the risk-mitigating effect of bonus cap and malus we found earlier becomes insignificant, once the earnings target is introduced: this evidence is consistent with our hypothesis (H3). In addition, bonus cap group was less likely to choose the lowest risk level (*Risklevel_1*) than the control group.

In addition, compared to their inherent risk preferences (choices in Task 0), we found that the introduction of an absolute earning target (£5m) significantly increases the number of participants who chose to risk up, across all three bonus groups. We fitted data to maximum-likelihood multinomial logit models with discrete dependent variables (i.e., Risk up, No change, Risk down) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(i\_task3 = 1)}{Pr(i\_task3 = 2)} = C_{10} + C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(i\_task3 = 3)}{Pr(i\_task3 = 2)} = C_{20} + C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where i_task3 denotes the change of participants' choices in Task 3 relative to those in Task 0, with i_task3=1, 2 and 3 corresponding to "risk down", "no change" and "risk up" respectively. BonusRegime (=1, 2 or 3) corresponds to proportional (the control), bonus cap and malus regimes. Inheritancechoice denotes the asset choices participants made in Task 0.

Table 11: Impact of bonus regimes on risking up in Task 3, relative to Task 0

	(1) i_task3	
risk_down		
Control	0.000	(.)
Bonus cap	-0.765	(0.494)
Malus	0.091	(0.452)
Inheritancechoice	0.852***	(0.229)
Constant	-3.572***	(0.956)
no_change		
Control	0.000	(.)
risk_up		
Control	0.000	(.)
Bonus cap	-0.239	(0.317)
Malus	-0.131	(0.326)
Inheritancechoice	-0.186	(0.131)
Constant	2.007***	(0.510)
Observations	392	
Pseudo R^2	0.053	

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11 shows that participants in bonus cap and malus groups were no less likely to risk up than the control group, and no more likely to risk down. There findings are consistent with our hypothesis (H3) that the risk-mitigating effect of bonus cap and malus weakens once bonus is made convex via the introduction of a high earnings target.

D. Impact of relative performance pay on the risk-mitigating effects of bonus cap and malus

In Study 2 Task 3A, participants were again presented the same 6-asset choice task as in Task 1 (see Table 2). In this task, the participants were informed that, in order to receive a bonus, their asset return has to exceed that of a competitor and that the competitor has chosen Asset 5.

H3': The risk-mitigating effect of bonus cap and malus weakens once a relative performance target is introduced.

Table A in the Annex shows that, compared with choices that participants made in Task 0, the introduction of this relative performance pay and the information that the competitor is taking a risky investment strategy increased the number of participants who chose the Assets 5 or 6 (*Risklevel_3*) across all three bonus groups.

We fitted data to maximum-likelihood multinomial logit models with discrete dependent variables (i.e., *Risklevel*=1, 2 or 3), controlling for participants' inherent risk preferences (i.e, their choices in Task 0) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(t3A\_risklevel = 1)}{Pr(t3A\_isklevel = 2)} = C_{10} + C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(t3A\_risklevel = 3)}{Pr(t3A\_isklevel = 2)} = C_{20} + \ C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where t3A_risklevel denotes the Risklevel (=1,2 or 3) participants chose in Task 3A, BonusRegime (=1, 2 or 3) corresponds to proportional (the control), bonus cap and malus regimes. Inheritancechoice denotes the asset choices participants made in Task 0.

Table 12: The impact of bonus regime on risk choice in the presence of relative performance benchmarking

	(1) t3A_risklevel	
Risklevel_1		
Control	0.000	(.)
Bonus cap	0.013	(0.567)
Malus	-0.363	(0.604)
Inheritancechoice	-0.422^{**}	(0.205)
Constant	0.765	(0.756)
Risklevel_2		
Control	0.000	(.)
Risklevel_3		
Control	0.000	(.)
Bonus cap	-0.093	(0.424)
Malus	0.193	(0.421)
Inheritancechoice	0.184	(0.162)
Constant	-0.202	(0.665)
Observations	173	
Pseudo R^2	0.032	

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12 shows that the risk mitigation effect of bonus cap and malus observed in Task 1 becomes statistically insignificant once the relative performance benchmarking is introduced. In addition, comparing with their inherent risk preferences (choices in Task 0), we found that the introduction of relative performance benchmarking significantly increases the number of participants who chose to risk up, across all three bonus groups. We fitted data to maximum-likelihood multinomial logit models with discrete dependent variables (i.e., Risk up, No change, Risk down) to test the statistical significance of the observed differences, and to identify the contribution from the different bonus regimes:

$$ln\frac{Pr(i\_task3A = 1)}{Pr(i\_task3A = 2)} = C_{10} + \ C_{11}(\text{BonusRegime} = 2) + C_{12}(\text{BonusRegime} = 3) + C_{13}\text{Inheritancechoice}$$

$$ln\frac{Pr(i\_task3A = 3)}{Pr(i\_task3A = 2)} = C_{20} + \ C_{21}(\text{BonusRegime} = 2) + C_{22}(\text{BonusRegime} = 3) + C_{23}\text{Inheritancechoice}$$

where i_task3A denotes the change of participants' choices from Task 0 to Task 3A, with i_task3A=1, 2 and 3 corresponding to “risk down”, “no change” and “risk up” in Task 3A relative to Task 0, respectively. BonusRegime (=1, 2 or 3) corresponds to proportional (the control), bonus cap and malus regimes. Inheritancechoice denotes the asset choices that participants made in Task 0.

Table 13: Impact of bonus regimes on risking up and down, with a relative performance metric

	(1)	
	i_task3A	
risk_down		
Control	0.000	(.)
Bonus cap	0.895	(0.788)
Malus	1.422*	(0.819)
Inheritancechoice	1.095***	(0.356)
Constant	-6.478***	(1.772)
no_change		
Control	0.000	(.)
risk_up		
Control	0.000	(.)
Bonus cap	0.381	(0.423)
Malus	0.772*	(0.441)
Inheritancechoice	-0.474**	(0.187)
Constant	2.102***	(0.738)
Observations	173	
Pseudo R^2	0.107	

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Compared with the control group, there is no longer definitive evidence of risk mitigation effect of bonus cap. For the malus group, we found that compared with the control group, participants in the malus group are marginally more likely to risk down and risk up. While the increasing likelihood of risk down can be attributed to the risk mitigation effect of malus as discussed in Task 1, increasing likelihood of risk up can be perhaps explained as stronger effect of herd behaviour of some participants. Overall, our results are consistent with our hypothesis (H3').

Based on the analysis above, we conclude that, while bonus cap and malus can mitigate risk-taking in principle, their risk-mitigating impact is easily undermined through an introduction of an absolute or a relative performance metric, both of which are commonly used by banks.

4. Analysis: Impact of bonus regime on effort

Next, we examine how the bonus regime affects project search effort. In both Study 1 and 2, participants were given real effort tasks. We use these tasks to examine the following hypothesis:

H4: Bonus cap and malus reduce incentives to engage in project search effort by reducing return from effort.

Study 1:

In Study 1 Task 2 and 4, participants were given a set of 30 assets which they had to examine by scrolling up and down the screen before choosing an asset to invest in (see Table 4). Whereas the participants were offered a bonus if they achieved a return greater than zero in Task 2, they were offered a bonus only if they achieved a return greater than £5 million in Task 4 (see Table 5).

Tasks 2 and 4 required cognitive effort in evaluating the risk-return characteristics of a large number of assets before making a decision. But there was no uncertainty in return from effort in these tasks, as all the probabilities and payoffs from all available investment opportunities were revealed to the participants. In our analysis, effort variable in these tasks received a binary classification: in particular, participants were deemed to have exerted effort (effort = 1) if they have chosen an asset from a set which dominates all other assets in the set, but otherwise deemed to have not exerted effort (effort = 0). We also examined the time taken to complete the task as an alternative measure of effort.

Table 14 summarises the results. It shows that, on average, 86.8% of the participants identified the correct set of assets in Tasks 2 and completed it in 222 seconds (3 minutes and 42 seconds), and 86.8% of them identified the correct set of assets in Task 4 and completed it in 175 seconds (2 minutes and 55 seconds). The bonus cap group had a lower percentage of participants identifying the correct set of assets than the control group both in Tasks 2 and 4. By contrast, the malus group had a higher percentage of participants identifying the correct set of assets in both tasks than the control. However, the difference between the control group and the bonus cap and the malus group in the likelihood of identifying the correct set of assets was not statistically significant. Similarly, the difference between the control group and the bonus cap and malus groups in the time taken to complete the tasks was not statistically significant. Thus, the results from Study 1 did not support our hypothesis (H4), either for bonus cap or malus.

Table 14: Study 1 Task 2 and Task 4, correct answers and average time taken to complete the tasks

	Average % correct	Average time taken (seconds)
Task 2		
Control	87.8%	211
Bonus cap	82.4%	227
Malus	90.1%	230
All	86.8%	222
Task 4		
Control	86.5%	172
Bonus cap	85.1%	189
Malus	88.7%	165
All	86.8%	175

<u>Study 2</u>:

The lack of evidence on the impact of bonus regime on effort in Study 1 could reflect the fact that the tasks were relatively easy, as evidenced by the fact that a high proportion of the participants could identify the correct set within a short period of time. In order to examine the hypothesis (H4) further, we set a harder project search effort task in Study 2. In this task (Task 2A), participants were given the first set of 6 assets shown in Table 4, but had to perform one 5-digit, 3 number summation or subtraction question correctly in order to reveal each of the additional 24 assets. Thus, participants could therefore view a maximum of 30 assets by performing all 24 calculations correctly, but they could also decide to stop doing the calculations at any point in time. After each computation, participants were informed whether their answer was correct or not.

Effort was measured by the number of questions *attempted*: this is the right measure in our setting, as this was the participants' choice variable. The number of correctly answered questions could be interpreted as output, or return on effort. Thus, correct answers per attempted questions could be interpreted as productivity.

Table 15 summarises our results. On average, the bonus cap group attempted 2.7 fewer questions than the control group, whereas the malus group attempted 1.2 questions fewer than the control group. Strikingly, 19.3% of the bonus cap group did not attempt any question, compared to 8.8% of the control group and 6.8% of the malus group (Figure 1). Figure 1 also shows that the bonus group was much more likely to give up very quickly after a few questions, compared to the control group.

Table 15: Study 2 Task 2 number of attempted and correctly answered questions, and correct answers per attempted questions

	Attempted		Correctly answered		Correct answers per attempted	
	Mean	Median	Mean	Median	Mean	Median
Control	11.5	9	9.3	8	79.8%	82.6%
Bonus cap	8.8	6	7.6	5	82.9%	88.2%
Malus	10.3	9	8.6	6	78.5%	85.7%
All	10.2	9	8.5	6	80.3%	85.7%

The OLS regression results on the number of math computations in Task 2 in Study 2 and number of correct answers are reported in Table 16. Our regression analysis of the number of attempted questions (first column, Table 16) shows that the negative coefficient on the bonus cap dummy was significant at 10% level. By contrast, the coefficient on the malus dummy was negative but not statistically significant. This result points to some evidence that the bonus cap could reduce effort, by reducing the expected return from effort. This evidence offers some support for our hypothesis (H4) for bonus cap only, but not for malus.

However, the bonus cap was not statistically significant for the number of correct computations, although the coefficient remained negative. This reflects the fact that, amongst those that decided to attempt at least one question, the bonus cap group had a marginally higher correct computation per attempted question than the control group, although the difference was not statistically significant. Other variables, such as age and sex were not statistically significant in any of these regressions.

Figure 1: Cumulative distribution, number of questions attempted by bonus group

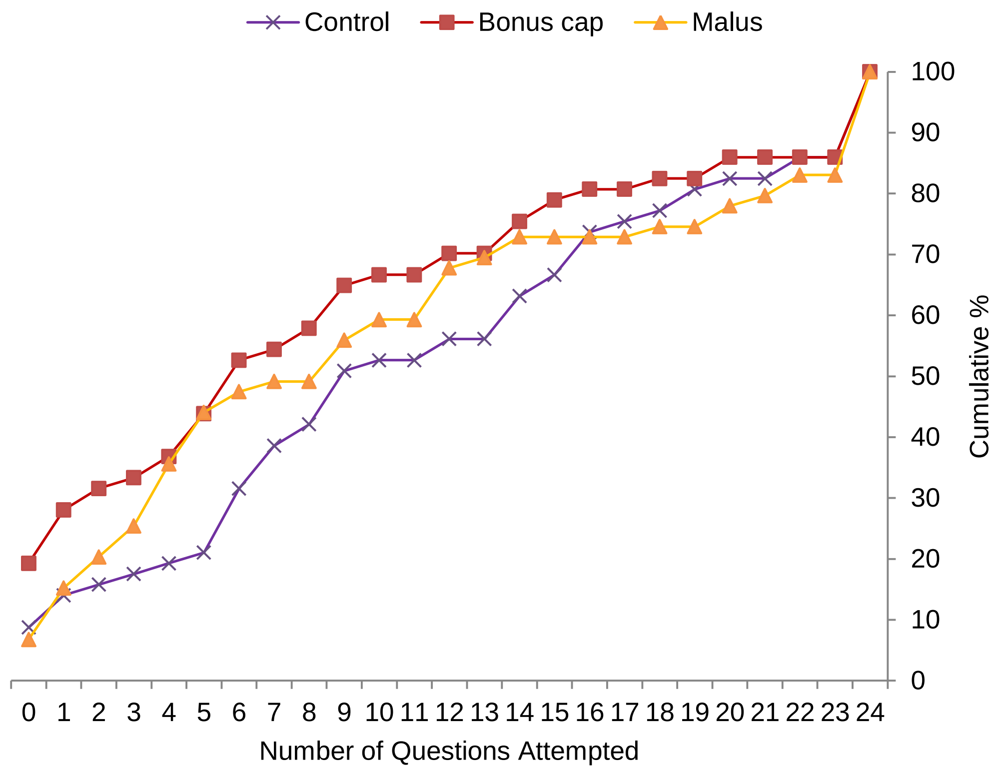

Table 16: Impact of bonus regimes on effort (number of computations) and productivity (number of correct answers)

	(1) No of computations attempted	(2) No of correct computations	(3) Share of correct computation per attempted
Bonus Cap	-2.664*	-1.630	0.019
	(1.497)	(1.315)	(0.045)
Malus	-1.281	-0.780	-0.011
	(1.484)	(1.304)	(0.043)
Math_tutorial	2.134***	2.093***	0.100***
	(0.640)	(0.562)	(0.020)
Constant	4.678**	2.581	0.473***
	(2.301)	(2.022)	(0.071)
Observations	173	173	153
Adjusted R^2	0.062	0.068	0.135

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5. Discussion and policy implications

Our findings suggest that, when individuals are under a bonus scheme which rewards them proportionally for positive investment returns but does not penalise them for negative returns (proportional bonus), they tend to take greater risks than they would with their own money (H1). The scenario in which participants are asked to invest their own money (Task 0) represents a 'frictionless' benchmark, in which there is no implicit principal-agent problem between the participant (agent) and other hypothetical 'stakeholders' (principal). Thus, our results can be interpreted as being consistent with the received wisdom that the proportional bonus scheme which offers rewards for positive returns but does not penalise for negative returns could encourage "excessive" risk-taking. We also find evidence that the imposition of bonus cap or malus conditions on such schemes can mitigate risk-taking incentives (H2). However, we also find that simple manipulations to the bonus structure, e.g. setting of an absolute or a relative performance target, are sufficient to undermine the risk-mitigating effects of bonus cap and malus (H3 and H3'). We also find some evidence that bonus cap might reduce project search effort (H4), but found no evidence that malus could reduce effort.

While it is not possible to draw direct inferences about the efficacy of the actual pay regulations based on our experimental study involving relatively small stakes, our results nevertheless offer several valuable insights into the potential weaknesses and limitations of these regulations. Importantly, our study shows how individuals' risk choices change when the bonus structure is altered while meeting specific constraints that mimic the existing pay regulations (bonus cap and malus). Thus, our findings point to the possibility that, as long as banks' remuneration committees, which represent their shareholders' interests, can freely vary the parameters that determine the incentive structures of bank executives, and the shareholders themselves want to encourage bank executives to take risks, pay regulations could only have weak impact in restraining risk taking.

This has two main implications. First, in order to monitor incentives facing bank executives, it may not be sufficient for regulators to check banks' compliance with the existing pay regulations, but it may be necessary to carefully examine the risk-taking incentives embedded in the entire pay structure. More specifically, regulators need to be tuned into the possibility that features such as absolute and relative performance targets could be used to fuel bank executives' risk-taking incentives even in the presence of pay regulations.[8] Second, in order to align the bankers' incentives with those of society, regulatory reforms aimed at eliminating distortions in the incentives of the bank shareholders – whose interests are ultimately mirrored in the bank executives' pay structure – could be more effective than regulating bankers' pay directly. The relevant regulatory reforms include those aimed at increasing shareholders' 'skin in the game' (e.g. via higher capital requirements and buffers) and ending too-big-to-fail (e.g. by improving resolvability of failed banks).

Our study also points to the possibility that, consistent with Hakenes and Schnabel's (2014) hypothesis, bonus cap could have the unwanted side effect of reducing the project search effort. Because bonus cap limits the potential reward from effort, it may be rational for individuals to 'shirk' when effort is costly. By contrast, we did not find any evidence that malus encourages shirking. This

[8] We also note that incentives could be manipulated by employment conditions other than bonus, e.g. promotions and sackings, over which regulators may have some, but not complete control. For example, in the United Kingdom, the Prudential Regulation Authority has powers to reject senior appointments at banks if they are deemed inappropriate.

result makes sense, as malus does not limit the potential reward from effort in the same way that bonus cap does. Thus, our result suggests that, while both bonus cap and malus share the same weaknesses, bonus cap is potentially more problematic: if the bonus cap actually reduces project search effort, it is not clear whether having this regulatory requirement helps or hinders more sensible risk choices.

6. Conclusions

The recognition that the bonus culture was a factor which led to the 2007-8 financial crisis led to the introduction of new regulations on bankers' pay across a number of jurisdictions. These new regulations were based on *ceteris paribus* reasoning: other things equal, the new regulatory requirements to cap bonuses or to penalise risk management failures through malus and clawback should lead to better alignment of banker executives' incentives. It is, however, more realistic to expect that banks will respond to these regulations by tweaking the pay structure, in order to retain bank executives' incentives to maximise shareholder returns (Thanassoulis and Tanaka, 2017). Thus, pay regulations are *robust* only if they can prevent excessive risk taking even when banks can adjust pay parameters that are under their controls. The empirical identification of the impact of pay regulations on incentives, however, is challenging to impossible due to the lack of data on individuals' decisions under different bonus regimes. In this context, our lab experiment provides a novel, alternative approach to improve our understanding of how these regulations might affect incentives.

Our study offers new evidence on how specific constraints imposed on bonus pay could influence risk taking and project search effort. First, consistent with the conventional wisdom, bonus that is proportional to positive investment returns but does not penalise for negative returns encourages risk taking. More specifically, we found that, under such a bonus regime, individuals take greater risks than they would with their own money, suggesting that such a regime could potentially encourage *excessive* risk taking. Second, we find that bonus cap and malus can mitigate this risk taking, *ceteris paribus*. Third, however, we also find that the risk-mitigating effects of bonus cap and malus can easily be undermined by the introduction of an absolute or a relative performance target. Finally, we also find some evidence that bonus cap might reduce project search effort, consistent with the theoretical prediction of Hakenes and Schnabel (2014), but we did not find evidence that malus encourages such 'shirking'.

Our findings are highly policy relevant. In particular, they suggest that the regulators' original diagnosis that the bonus culture was a factor that led to the 2007-8 financial crisis may well have been right. Nevertheless, our findings raise questions over the efficacy of regulating bankers' pay when banks' shareholders want to encourage their executives to take greater risks than what taxpayers would prefer. Such a divergence in interests between banks' shareholders and taxpayers is likely to remain as long as the shareholders do not bear the full cost of banks' risk-taking due to the implicit and explicit government guarantees on bank liabilities provided through mispriced deposit insurance and the inability of the authorities to fully rule out the possibility of a bailout using public funds. Thus, the first best solution for aligning bankers' incentives would be to address shareholders' incentives directly, for example through reforms that make the resolution of large, systemic banks more credible. If, however, it is not feasible to fully eliminate the distortions in shareholders' incentives, then regulators need to not only monitor compliance with pay regulations, but also examine bank executives' incentive pay more holistically in order to identify features that could potentially encourage excessive risk taking, even in the presence of bonus regulations.

Annex:

A1. Probability tutorial questions

This part is intended to help remind you of the basic probability theory that you might find useful for completing Part 3 later. To warm you up for questions in Part 3, you will be asked three questions involving probability.

Question 1: If you flip a fair coin three times, what is the probability of getting three heads?

- a) 1/8 = 12.5%
- b) 1/6 = 16.7%
- c) 1/4 = 25%
- d) 1/2 = 50%

Correct answer: a) The probability of three heads = 50% x 50% x 50% = 12.5%

Question 2: Suppose you enter a gamble involving flipping a fair coin multiple times. You will win £10 every time you see heads, and lose £5 every time you see tails. If you flip the coin 10,000 times, how much do you expect to win or lose every time you flip the coin on average?

- a) Lose £2.50.
- b) Neither lose nor gain any money (£0).
- c) Win £2.50.
- d) Win £5.00.

Correct answer: c). On average, you should expect to make £2.50 each time you flip the coin: 50%*£10 + 50%*(-£5) = £2.50.

Question 3: Suppose you are an investment manager and you are evaluating an opportunity to invest £100 million. The project is expected to succeed with 70% chance and yield a return of £10 million, but the project could fail with 30% chance and result in a £5 million loss. You will be paid a bonus proportional to the return on the investment at a rate of £1,000 per million return on your investment, but you will be paid no bonus if the project results in a loss.

What is the expected return on this investment?

- a) £5 million
- b) £5.5 million
- c) £7 million
- d) £8.5 million

What is your expected bonus?

- e) £5,000
- f) £5,500
- g) £7,000
- h) £8,500

Correct answers: The correct answer for the expected return on this investment is b) £5.5 million (70 %*£10 million + 30%*(-£5 million) = £5.5 million). The correct answer for your expected bonus is g) £7,000 (70 %*£10,000 + 30%*(£0) = £7,000).

A2. Table A: Asset choices in 6-asset risk choice tasks

i) Frequency

	Study 1								
	Control			Bonus cap			Malus		
	Task 0	Task 1	Task 3	Task 0	Task 1	Task 3	Task 0	Task 1	Task 3
Asset 1	1	2	3	8	3	1	1	5	3
Asset 2	12	10	5	15	12	2	9	14	6
Asset 3	20	16	10	12	20	3	14	19	9
Asset 4	41	38	8	34	38	7	46	32	12
Asset 5	0	2	37	2	0	55	1	1	39
Asset 6	0	6	11	3	1	6	0	0	2
Total	74	74	74	74	74	74	71	71	71

	Study 2											
	Control				Bonus cap				Malus			
	Task 0	Task 1	Task 3	Task 3A	Task 0	Task 1	Task 3	Task 3A	Task 0	Task 1	Task 3	Task 3A
Asset 1	6	4	4	1	3	2	2	4	3	6	2	0
Asset 2	6	5	3	4	6	2	2	2	7	10	3	1
Asset 3	6	9	5	5	9	21	8	4	13	14	2	6
Asset 4	32	27	7	18	34	28	15	19	34	28	12	18
Asset 5	3	4	27	17	2	0	25	17	0	0	37	24
Asset 6	4	8	11	12	3	4	5	11	2	1	3	10
Total	57	57	57	57	57	57	57	57	59	59	59	59

ii) As percentage of total within each bonus groups

	Study 1								
	Control			Bonus cap			Malus		
	Task 0	Task 1	Task 3	Task 0	Task 1	Task 3	Task 0	Task 1	Task 3
Asset 1	1.4	2.7	4.1	10.8	4.1	1.4	1.4	7.0	4.2
Asset 2	16.2	13.5	6.8	20.3	16.2	2.7	12.7	19.7	8.5
Asset 3	27.0	21.6	13.5	16.2	27.0	4.1	19.7	26.8	12.7
Asset 4	55.4	51.4	10.8	45.9	51.4	9.5	64.8	45.1	16.9
Asset 5	0.0	2.7	50.0	2.7	0.0	74.3	1.4	1.4	54.9
Asset 6	0.0	8.1	14.9	4.1	1.4	8.1	0.0	0.0	2.8

	Study 2											
	Control				Bonus cap				Malus			
	Task 0	Task 1	Task 3	Task 3A	Task 0	Task 1	Task 3	Task 3A	Task 0	Task 1	Task 3	Task 3A
Asset 1	10.5	7.0	7.0	1.8	5.3	3.5	3.5	7.0	5.1	10.2	3.4	0.0
Asset 2	10.5	8.8	5.3	7.0	10.5	3.5	3.5	3.5	11.9	16.9	5.1	1.7
Asset 3	10.5	15.8	8.8	8.8	15.8	36.8	14.0	7.0	22.0	23.7	3.4	10.2
Asset 4	56.1	47.4	12.3	31.6	59.6	49.1	26.3	33.3	57.6	47.5	20.3	30.5
Asset 5	5.3	7.0	47.4	29.8	3.5	0.0	43.9	29.8	0.0	0.0	62.7	40.7
Asset 6	7.0	14.0	19.3	21.1	5.3	7.0	8.8	19.3	3.4	1.7	5.1	16.9

References

Andersson, O., Holm, H.J., Tyran, J-R, and Wengstrom, E. (2013), "Risking Other People's Money". *CEPR Discussion Paper*, DP 9543.

Brüggen, A. and Strobel, M. (2007), "Real effort versus chosen effort in experiments". *Economic Letters*, 96, 232-236.

Cole, S., Kanz, M., and Klapper, L. (2015), "Incentivising calculated risk-taking: evidence from an experiment with commercial bank loan officers". *Journal of Finance*, No. 70 (2), 537-575.

Dohmen, T. and A. Falk (2011), "Performance pay and multidimensional sorting: Productivity, preferences, and gender". *American Economic Review,* 101 (2), 556-590.

Edmans, A., Gabaix, X. and Jenter, D. (2107), "Executive compensation: a survey of theory and evidence". Mimeo.

Edmans, A. and Liu, Q. (2011), "Inside debt". *Review of Finance*, 15, 75-102.

Eriksson, T., Poulsen, A., & Villeval, M. C. (2009). Feedback and incentives: Experimental evidence. *Labour Economics*, 16(6), 679-688.

Financial Stability Forum (2009). *FSF Principles for Sound Compensation Practices*. Available from: http://www.fsb.org/2009/04/principles-for-sound-compensation-practices-2/

Hakenes, H. and Schnabel, I. (2014), "Bank bonuses and bailouts". *Journal of Money, Credit and Banking*, 46, 259-288.

Hartmann, F. and Slapnicar, S. (2015), "An experimental study of the effects of negative, capped and deferred bonus on risk taking in a multi-period setting". *Journal of Management and Governance*, vol. 19, 875-896.

Kühberger, A. Schulte-Mecklenbeck, M. and Perner, J. (2002). "Framing decisions: hypothetical and real". *Organizational Behavior and Human Decision Processes*, vol. 82, 1162-1175.

Levitt, S. D. and Neckermann, S. (2014). "What field experiments have and have not taught us about managing workers". *Oxford Review of Economic Policy*, Vol.30, No. 4, 639-657.

Niederle, M. and Vesterlund, L. (2007). "Do women shy away from competition? Do men compete too much?". *Quarterly Journal of Economics,* 122 (3), 1067–1101.

Thanassoulis, T., and Tanaka, M. (2017), "Optimal pay regulation for too-big-to-fail banks", *Journal of Financial Intermediation*, https://doi.org/10.1016/j.jfi.2017.03.001.