Essays on Counter Strategies Against Mistakes in Information Processing with Various Applications

Thesis by:

Moritz Mosenhauer, M.Sc.

Submitted in fulfillment of the
requirements for the Degree of:
Doctor of Philosophy in Economics

Adam Smith Business School, College of Social Science
University of Glasgow
June, 2020

This page is intentionally left blank.

Abstract

This thesis consists of three chapters, dealing with mistakes in information processing and how smart information design may improve upon it.

Chapter 1 studies a firm's organisational responses if its agents misevaluate information. If a manager overreacts to unusual events, it may be desirable for the firm to adopt the management practice *management-by-exception*. I develop a theoretical framework to study this technique and derive conditions under parsimonious assumptions for when it should be adopted. Moreover, I show how further assumptions can refine the model's predictions, establishing a direct link between the manager's over-responsiveness and organisational rigidity. The strategy is implemented by controlling the information that the manager receives. In fact, in the absence of information transmission and processing costs, it may be optimal to not send inherently valuable signals concerning the economy's state to the manager.

Chapter 2 investigates tools to counter excessive stock trading and increase profits for private households participating in the stock market. Creating a stylised hold or trade-scenario in a computer laboratory experiment, I find that by solely changing the information the participants receive, trading activity can be reduced by roughly 30%, increasing trading profits by more than 0.55 percentage points on monthly net returns. In particular, I consider two information treatments. First, I provide the participants with additional information by giving detailed feedback on their actions and outcomes at every turn. Second, when considering whether to hold a given stock or trade it for another one, I restrict participants' information on the recent performance of their allocated stock. Both interventions lead to significant changes in behaviour. Additionally, the 2×2 experimental design reveals that the effects stack.

Chapter 3 deals with growth diagnostics. Growth diagnostics is an influential policy framework that, in second-best settings, has been used to identify the priorities of policy reform in different countries. With limited information about the nature of interaction between different second-best distortions, mistakes (situations where realized social welfare losses overwhelm any intended gains) could occur in the implementation of growth diagnostics. Even allowing for the possibility of mistakes, would an adaptive implementation of growth diagnostics converge to a socially optimal outcome? This paper sets out the conditions under which such convergence occurs. A number of different historical examples are discussed to illustrate how such a process could play out in practice.

Contents

List of Tables

List of Figures

Acknowledgements

I thank Natacha Estèves for being the greatest thing in my life.

Throughout all stages of my life and studies, my parents Helga and Goerg Mosenhauer have given me constant support and love.

Special gratitude belongs to my supervisors Sayantan Ghosal and Yiannis Vailakis. My work has profited in innumerable ways from their encouragement, critique, insight and backing. It has been an honour and privilege to have been their student.

I further would like to single out Theodore Koutmeridis and Bram Driesen from the University of Glasgow's Economics Department. Both have acted as great mentors and always provided an open ear for me. I am deeply thankful for their selfless commitments. Also, I thank Florian Englmaier for his expert advice on my job market paper and writing way too many recommendation letters.

Unfortunately, my work has benefited, directly or indirectly, from too many too count or mention here. To name just a few, I thank Michele Battisti, Tanya Wilson and Souvik Datta for their input on my job market preparation; Claudio Mezzetti, Mark Dean and Andrew Clausen for their feedback on my work as well as Lorna Wilson and Sandra Dopico-Ardao for their friendly and efficient help.

Last but not least, I would like to thank all my Ph.D. colleagues at the University of Glasgow. Continuously, they have sacrificed hours to help conduct my experiments, torn my research into shreds in our gruelling reading groups and helped on missing links in proofs and notation lacking rigour. I am thankful for their input and making my time in Glasgow unforgettable.

Affidavit

I declare that, except where explicit reference is made to the contribution of others, this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Moritz Mosenhauer

Innsbruck, June 29, 2020.

Chapter 1

Salience and Management-by-Exception

1.1 Introduction

Consider a manager that needs to correctly set priorities for allocating scarce resources within a firm across its sectors, say, production and marketing, in order to adapt the work processes better to changing circumstances. If the manager is perfectly rational, more information, such as knowledge of demand shocks in the economy or changes in the demographics of targeted buyers, will always improve this judgement. However, if the manager systematically misevaluates information, it may also introduce a bias and thereby drive her away from the correct choice. In this case, the manager's decision-making may be improved by providing less information.

I show that a firm's system of management practices may take up the role of aiding the manager by acting as a smart pass-through for information. In particular, I provide a behavioural foundation for a firm's choice to only brief the manager on sufficiently unusual events, *i.e.* adapt *management-by-exception* (see Section 1.2.2 for a verbal definition and a discussion of the concept), when the manager overreacts to *salient* information (see Section 1.2.1 for a conceptual discussion and empirical evidence on behavioural biases in managerial decision-making).

This paper deals with firms that do not face conflicts of interest: two individuals, a manager (she) and an owner (he), attempt to maximise a common objective. The manager directly affects the firm's profit. In particular, she decides how to use scarce resources in order to adapt the firm's work processes to a changing productive environment. Information regarding the state of the economy, *per se*, is a valuable signal in doing so. However, the manager has a tendency to misevaluate this information when incorporating it into her decision-making by overemphasising salient events. This may lead to a misallocation of resources and, thus, a failure to maximise profits.

In this setting, it is impossible to achieve the first-best outcome implied by perfectly informed and non-biased managerial decision-making. But the owner may still pursue a second-best solution by deciding whether or not and, more precisely, to what degree to implement the management technique *management-by-exception*. He does so by merely varying the information available to the manager on the economy's state so that she is only briefed if exceptional occurrences arise. Restricting access to any information concerning comparatively usual events, the owner may eliminate the manager's psychological bias through evoking adherence to an *ex ante* optimal plan, but he will also decrease the firm's adaptability. On the other hand, for comparatively exceptional events the owner may wish to forfeit the institutional rigidity and inform the manager of the unusual occurrences, however, at the cost of reintroducing the bias.

Investigating the owner's trade-off between either a better informed or a less biased decision-making by of the manager, as captured by the implementation of *management-by-exception*, shall be the main focus of this article. Under parsimonious assumptions, I first identify manager- and production-specific characteristics driving the adoption of this management practice (see Proposition 2). More fine-grained predictions can further be made if the state of the economy is assumed to be standard normally distributed: then, for all possible productive settings there exists a uniquely optimal way of defining what should constitute an 'exception' in guiding managerial attention (see Proposition 3).

With this article, I make three major contributions. I formulate a possible organisational coping strategy when workers within the firm are affected by psychological biases in valueing their options. My research thus aligns with the seminal work of Simon (1945), who places aspects of bounded rationality at the centre of the debate on how to optimally design organisations. While a host of other theoretical studies work in this tradition, they have so far largely been restricted to information transmission constraints as the sole imperfection in the agents' decision-making (see Dessein and Prat, 2016, for a review). In going beyond this, my research answers recent mounting evidence suggesting that behaviourally biased valuation by top managers is, in fact, a prevalent real-life phenomenon.

Moreover, I present a novel model providing a theoretical foundation for a firm's choice of whether to adopt *management-by-exception*. Following Bertrand and Schoar (2003), there has been a surge of interest in the financial and economic literature in explaining the adoption choice of management styles and techniques by organisations and their effectiveness in practice (*e.g.* Bloom and Reenen, 2007; Benmelech and Frydman, 2015; Bloom et al., 2017). Long preceding this trend, extensive research efforts have been made in the management literature pursuing the same objective. Based on Bass (1985)'s extension of the transactional-transformational leadership paradigm, Bass designed the *Multifactor Leadership Questionnaire* survey including tokens specifically designed to test for the implementation of *management-by-exception*. With more than 80 studies conducted using the survey-tool, a number of meta-analyses have attempted to empirically summarise the findings.

While these studies successfully uncover some stable relationships between management practices and measures of effectiveness, *management-by-exception* remains a 'problem child'. (Judge and Piccolo, 2004, p. 755) state that *'management by exception (active and passive) was inconsistently related to the criteria'.* Similarly, (Lowe et al., 1996, p. 416) note that *'Management-by-Exception [was] inconsistent in [its] relationships with effectiveness across studies. Some research evidenced positive associations, while other findings showed a negative association'.*

While there have been previous studies that capture how *management-by-exception* is implemented in a firm, these do not allow a meaningful analysis of whether or not a firm should adopt it. In Athey et al. (1994), the adoption of the technique is assumed, while in Garicano (2000) and Beggs (2001) it is inextricably related with the presence of a hierarchy. The adoption of the technique is thus given apart from the unrealistic case of a firm with only workers and no managers. Also, all previous models feature some explicit or implicit costs of information transmission. I show that firms may benefit from implementing *management-by-exception* even in the absence of any such costs. By considering a different set of underlying drivers, I derive a set of new comparative statics which not only provide new predictions but also help illuminate previous findings (see Section 1.4).

Lastly, even though virtually all of the literature on *management-by-exception* agrees on the importance of establishing a plan that allows identifying exceptions, there only exists vague guidance on how such a system should be set up and calibrated. For example, *management-by-exception* has in recent years received interest towards applications in the automation of unmanned aircraft. Such aircraft may perform all necessary functions automatically while skilled personnel are ready on the ground to remotely take control over the machine if unusual circumstances arise. However, when practically implementing the system it has been observed that 'defining the basis for switching levels of automation support to the human remains difficult' (Dekker and Woods, 1999, p.

88). Experience-based rule-of-thumb and intuitive graphical analyses appear to be common methods in practice. To the best of my knowledge, there does not exist a rigorous theoretical foundation to answer a central question: *"What is an exception?"* (Dekker and Woods, 1999, p. 88) In Section 1.5, I show that this question can, in fact, be answered precisely in the given setting if sufficiently strong assumptions are placed on the distribution of the task-level shocks.

1.1.1 Related Literature

The article should be read as complementary to the recent literature in team theory and organizational economics featuring imperfect agents. A host of studies using a similar functional setting that favours adaptation to local shocks have pointed out the desirability of institutional rigidities. Dessein et al. (2016) find that if there are explicit costs of coordination across tasks, it may be optimal to fix the firm's operations for a large number of its activities and focus attention on a small number of core competencies. In Powell (2015), static bureaucratic rules prevent managers from engaging in costly influence-activities in an attempt to shape the organisation's outcomes towards their personal goals. However, in both of these models the communication channels and governance structures, respectively, cannot respond to the realisation of the shock vector, precluding any chance of flexibility in the organisational framework and to *manage-by-exception* in particular.

Interestingly, the main trade-off in this paper — better information versus a less-biased decision-making — is essentially identical to that of Dessein (2002), while it is arrived at and resolved from two different angles. In Dessein (2002), a principal holds decision rights and decides whether to delegate or simply gather information from an agent that strategically communicates. In my paper, delegation is fixed to a manager who is naive towards the bias in her decision-making and therefore behaves non-strategically. It is then the principal who strategically chooses the amount of information supplied to the manager in order to favourably affect the outcome.

In spirit, this paper is closest to Dessein and Santos (2016) in dealing with optimal organisational responses to managers who are biased in their valuation of the economy's state towards certain tasks. Methodologically, however, this article can best be understood as a team-theoretical, bounded-rationality take on Brocas and Carrillo (2007), where the decision of information transmission versus information restriction is made with respect to the size of the local shocks instead of iterative revealing of noisy signals. As an interesting sidenote, this article shows that it is possible to relax the strict assumption of information being public by assuming both the distribution of shocks as well as the information restriction device to be symmetric.

The literature on Bayesian Persuasion and Information Design studies a similar problem as this paper: a set of principals attempt to affect a set of agents' choices by selecting and committing to an information structure which maps possible states of the world into signals. Using the terminology of Bergemann and Morris (2016), this paper's setup most closely resembles the case of Bayesian Persuasion with an uninformed receiver, since there exist one agent and one principal where the principal has an informational advantage over the agent[1].

Due to some key differences, however, my work should be viewed as distinct from this literature. Instead of dealing with players with inherently conflicting interests, the principal in my case attempts to rectify decision-making mistakes stemming from the agent's distorted evaluation of salient states of the economy. In this scenario, it is natural to adopt a continuous state space, since the agent could otherwise perceive states which she knows cannot exist. The two most common solution procedures used in the literature, concavification (Kamenica and Gentzkow, 2011) and linear programming over Bayes-Correlated Equilibria (Bergemann and Morris, 2019), do not apply when there are infinitely many states. While Gentzkow and Kamenica (2016) and Ivanov (2015) have developed procedures allowing for infinite state spaces, these place some crucial restrictions on the preferences of the principal, namely that the economy's states do not enter them directly. Since in this article's setup the principal does not only care what action the agent takes but rather whether the agent's action constitutes a mistake for a certain state of the world, these approaches should be deemed unfitting for the problem under study.

More generally speaking, my findings corroborate articles that are driven by explicit or implicit costs of information transmission or processing such as Sah and Stiglitz (1986), de Clippel et al. (2017) and Calvó-Armengol et al. (2015). While this paper takes a reduced-form approach in modelling the agents' mistakes, some articles take a step towards endogenising decision imperfections by letting agents react optimally to limited information processing capabilities in the sense of Sims (2003)'s 'rational inattention' (see e.g. Dessein et al., 2016; Dessein and Santos, 2016).

Although the saliency bias enters this paper's model exogenously, its results itself can be interpreted as endogenising decision mistakes. In particular, the mechanism

[1]As noted in Bergemann and Morris (2019), although in this paper the principal commits to an information structure before observing the signal, it suffices that the final signal realisation can be conditioned on the true state of the world.

can be understood as a device for efficient attention allocation by negotiating *top-down* and *bottom-up* drivers in guiding focus (see Wolfe et al., 2003), especially if an intra-personal view on the model is adopted (see footnote 4). The model's framework may thus be utilised to provide further structure to the *tunnelling*-mechanism described in Mullainathan and Shafir (2013) or even consumers reacting to *extraordinary events* in Reis (2004).

1.2 Discussion of related concepts

1.2.1 Salience and Behaviourally Biased Managers

Previous research shows that the behaviour of lower-level workers is partly driven by behavioural, non-monetary motivations (see for example Blanco et al. (2017) or Danilov and Sliwka (2017)). There is also ample evidence, however, suggesting that the decision-making of top managers is influenced by psychological factors. Malmendier and Tate (2005) show that measures of overconfidence predict investment decisions of Forbes 500 managers. Malmendier et al. (2011) and Bernile et al. (2017) find that even formative events long before the managers' careers start, such as exposure to economic or natural disasters, shape their strategies of corporate financing and risk-taking.

Some of these behavioural biases appear to work through how managers incorporate information into their decision-making. In an early, impactful article, Ocasio (1997) cites the 'saliency of issues and answers' (p. 195) as one of six main mechanisms governing management choices. The role of saliency in individual decision-making is a well-established psychological phenomenon and has received considerable attention from economics. It posits that people do not treat all information equally, but instead some attributes may involuntarily make some pieces of information 'stand out' and cause them to subsequently receive a disproportionate weight in their decision-making (see Taylor and Fiske (1978) for a review in psychology and DellaVigna (2009) for one in economics).

A series of articles corroborate this view. Dittmar and Duchin (2016) show that personal experience has a stronger influence on a manager's later behaviour if it is gathered during a salient period in their career, while Gennaioli et al. (2016) attest that top CFOs overvalue recent observations when predicting earnings. Englmaier et al. (2017) offer evidence from a randomised controlled trial where they treat a number of managers with altered information designs regarding the incentive schemes while keeping actual compensation unchanged. They, too, argue that differences in behaviour stem from increased salience of certain dimensions of the productive process.

While Kőszegi and Szeidl (2013) formulate a model with a similar focus, I will closely follow Bordalo et al. (2013) in functionally capturing the manager's systematic bias (see Section 1.3.2 for the mathematical formulation). In Bordalo et al. (2013)'s setup, it is those attributes that differ most strongly from what is considered as 'usual' that gain salience. Results from Barber and Odean (2008) seemingly underpin this view. Studying individual stock traders, they find that 'stocks experiencing high abnormal trading volume (...) and stocks with extreme one-day returns' (Barber and Odean, 2008, p. 785) are bought disproportionately much.

1.2.2 Management Practices

Management-by-exception embodies a systematic decision-making procedure for evaluating and reacting to challenges encountered by an organisation. Bass (1990) provides a concise definition of the concept, stating a leader who manages by exception

> *watches and searches for deviations from rules and standards, takes corrective action (p. 22)*

The Business Dictionary (2018) elaborates further on this, defining *management-by-exception* as a

> *practice whereby only the information that indicates a significant deviation of actual results from the budgeted or planned results is brought to the management's notice. Its objective is to facilitate management's focus on really important tactical and strategic tasks.*

The management literature usually distinguishes between *active management-by-exception* and *passive management-by-exception.* The former is conceptually close to *laissez-faire leadership* where the leader prefers to remain inactive unless exceptional problems arise. This paper studies the former, as the manager will always take action within the model. As the crucial hallmark of the management technique I identify its information management, namely that the manager is only briefed on exceptional occurrences. Therefore, the benchmark to which the adoption of *management-by-exception* will be compared throughout the paper is the scenario where the manager is fully informed at all times.

The concept has long been subject to academic debate, with an early treatment of the general idea dating back to Towne (1886). It has gained substantial prominence in the management literature since Bass (1990) included it as a core characteristic of leadership in his extension of the transactional-transformational paradigm introduced by Burns (1978). In fact, (Patterson et al., 1995, p. 3) argue that "over the past decade there has probably not been a more dominant paradigm in leadership thought".

Management-by-exception also enjoys widespread use in practice. Mackintosh (1978) reports that by the 1970s the technique already enjoyed widespread use in a large variety of fields. Nowadays, it is one of seven core principles of the *PRINCE2* system, one of the world's most prevalent structured project management methods with over 1.4 million certified graduates[2] from a diverse array of countries and economic sectors (Axelos, 2016).

While the intended application of this practice was business management, *management-by-exception* has recently received further interest from air traffic control and specifically the automation of unmanned aircraft. Unmanned aircraft are already widely used in military operations (Hottman and Sortland, 2006) and is 'on the verge of taking flight alongside manned aircraft in the National Airspace System' (Liu et al., 2013, p. 424) in the United States of America. Such aircraft may perform all necessary functions automatically while skilled personnel are ready on the ground to remotely take control over the machine if unusual circumstances arise. There appears to be a consensus in the relevant literature that a major bottleneck for a successful introduction of these systems lies in 'defining the basis for switching levels of automation support to the human' (Dekker and Woods, 1999, p. 88) and that 'an appropriate level of automation is critical to the safety and performance characteristics of [unmanned aircraft systems] design.' (Liu et al., 2013, p. 425)

Analogous to this, this article characterises *management-by-exception* as an interplay between *ex ante* (before observing the current state of the world) and *ex post* (after observing the current state of the world) modes of decision-making. Following Dessein et al. (2016), I borrow vocabulary from March and Simon (1958)[3] in identifying these two modes of decision-making as *management-by-plan* and *management-by-feedback*, respectively. Investigating the optimal, exception-based switching point between plan and feedback management shall in fact be the main interest of this study.

[2]Database freely accessible at `https://www.axelos.com/successful-candidates-register`, accessed 19th of July, 2018

[3](March and Simon, 1958, p.182) remark: *'We may label coordination based on pre-established schedules coordination by plan, and coordination that involves transmission of new information coordination by feedback.'*

1.3 The Model

I posit a one-shot sequential-move principal-agent model,[4] which I solve for a second-best solution by backward induction. Decision rights over the firm's operations are fixedly delegated to the agent (henceforth called the manager, *she*) who is the only one that may directly affect the firm's profits. However, she naively and irreversibly makes mistakes in incorporating information upon observing the state of the economy, which introduces a systematic bias into her decision-making and causes her to potentially fail in maximising the firm's profits. The principal (henceforth called the owner, *he*), can influence the actions of the manager by controlling her access to information. In doing so, he faces a trade-off between a better-informed manager and a less-biased decision-making. Figure 1.1 shows the model's timeline[5].

Figure 1.1: Timeline of the model.

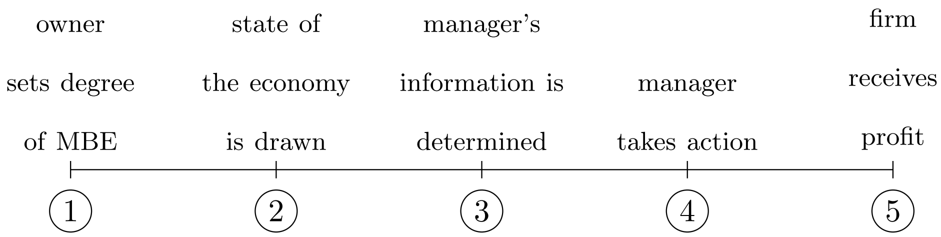

I follow much of the relevant literature in assuming a tracking-cost framework as the functional setting (Dessein, 2002; Rantakari, 2008; Calvó-Armengol et al., 2015; Powell, 2015; Dessein et al., 2016). Moreover, I choose a team-theoretic setting where all individuals in the firm, the manager and the owner, attempt to maximise the firm's expected profits, although the manager's propensity to misevaluate signals may lead to a failure to do so. The firm earns profits π across two tasks, or sectors, indexed with s (*e.g.* production and marketing):

$$\pi \equiv \sum_{s=1}^{2} \left[K_s - \beta_s \cdot |\theta_s - a_s(m(\omega))|^{\lambda} \right] \tag{1.1}$$

[4] I opt to model the two decision-makers within the organisation, the manager and the owner, as two separate individuals. This is, however, simply due to expositional convenience. Apart from their different action spaces in the model, the manager's role is to unconsciously be biased by psychological factors while the owner's role is to take this bias into account and take precautions in order to avoid adverse consequences arising from it. In the spirit of Thaler and Shefrin (1981), the inter-personal problem could equivalently be understood as an intra-personal problem of a single individual with two selves. Although the manager may be inherently unable to reverse her own tendency to misevaluate, she may be aware of it and put mechanisms into place in order to commit her future self to a restricted set of information.

[5] It should be noted that actions are only taken in period 1 and 4. I opt for describing the model in extensive form as it clarifies key aspects of it, e.g. the owner commits to a choice before the state of the world is realised as well as an inherent difference between the realisation of the state of the economy and the manager's information.

where K_s is the fix revenue for task s which, without loss of generality, is normalised to 0 for both s; θ_s are real numbers which are drawn from a single distribution with density function $g(\theta_s)$; β_s are strictly positive real numbers assigning *ex-ante* weights to the two tasks; λ is a strictly positive real number governing the curvature of the costs across both tasks; and $m(\cdot)$ is the choice of the manager made on the basis of information ω, which translates into firm-level outcomes via a_1 and a_2.

The firm incurs costs according to the random task-level shocks θ_s. Throughout the article I will maintain the following assumption:

Assumption 1. *Let the density function $g(\theta_s)$ be such that for any s*

(i) $E[\theta_s]$ exists

(ii) $g(\theta_s)$ is continuous

(iii) $g(\theta_s)$ is symmetric around 0.

It is the manager's role to shape the firm-level outcomes a_1 and a_2 by appropriately choosing $m(\omega)$ in order to avoid such costs. Thus, the firm's objective is to adapt as best as possible to changes in the productive environment.

1.3.1 The Owner

The owner cannot directly influence the profits of the firm. He can, however, affect the decision-making of the manager by deciding whether or not and, more precisely, to what degree to implement the management technique *management-by-exception* (see Section 1.2.2 for a discussion). He does so by purposefully controlling the amount of information available to the manager, denoted by $\omega \in \Omega$. At all times, ω contains at least the following information: *(a)* the profit function π, including full knowledge of the parameters β_1, β_2 and λ, *(b)* all available information on the density function of the task-level shocks $g(\theta_s)$ and *(c)* the structure of the problem.[6] This fundamental set of information is denoted with $\underline{\omega}$.

The manager will then either remain with her fundamental knowledge $\underline{\omega}$ or she will additionally receive full information on the task-level shocks θ_1 and θ_2, resulting in the final knowledge denoted $\overline{\omega} := \{\underline{\omega}, \theta_1, \theta_2\}$. This hinges on the owner's choice of implementing *management-by-exception.* If he does not adopt the practice, the manager will always receive full information on the state of the economy. If, on the other hand,

[6]In the current setup, it makes no difference to the model's results whether and to what extent the manager can observe the owner's action. I discuss this point after introducing the manager's action space in Section 1.3.3.

he does, the manager will only receive additional information once exceptional circumstances arise. I model the owner's choice via a real number $z \in [0, \infty)$. This choice affects the information available to the manager in the following way.

$$\omega = \begin{cases} \underline{\omega} & \text{if } |\theta_1| \wedge |\theta_2| < z \\ \overline{\omega} & \text{if } |\theta_1| \vee |\theta_2| > z \end{cases}$$

If the realisation for at least one of the realisations of the task-level shocks θ_1 or θ_2 deviates by more than z from their expected value (which by Assumption 1 is normalised to 0), the manager will receive full information on both values. In this case, events are sufficiently extreme to be considered an 'exception'. If, on the other hand, both realisations of the task-level shocks are sufficiently moderate so that neither of them deviate by more than z from their expected value, the manager receives no information on the state of the economy but instead makes her decisions on the basis of her fundamental knowledge $\underline{\omega}$. In this case, the firm concludes that no 'exception' has occurred.

Given this characterisation, I identify the implementation of *management-by-exception* within the model as well as different degrees of *management-by-exception* in the following way.

Definition 1. *The firm implements management-by-exception if the owner chooses any strictly positive value for z. Different values of z correspond to different degrees of management-by-exception being implemented, where lower values of z correspond to a more sensitive organisation and higher values of z to a more rigid organisation.*

1.3.2 The Manager

The manager attempts to adapt the firm's operations to the productive environment. In order to add economic relevance to the problem, I implicitly assume that the firm has limited resources at its disposal so that the operations cannot be perfectly adapted to the state of the economy. The manager must therefore decide how to prioritise the firm's tasks when allocating these resources. I formalise this decision as a binary choice based on the manager's available information.

$$m : \Omega \to \{1, 2\} \tag{1.2}$$

This decision then translates into the firm-level outcomes a_s in the following way.

$$\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \Big(m(\omega)\Big) := \begin{cases} \begin{pmatrix} \theta_1 \\ E[\theta] \end{pmatrix} & \text{if } m(\omega) = 1 \\ \begin{pmatrix} E[\theta] \\ \theta_2 \end{pmatrix} & \text{if } m(\omega) = 2 \end{cases}$$

Without intervention by the manager, work for both tasks is carried out doing 'business-as-usual' by setting $a_1 = a_2 = E[\theta]$. The manager then improves on this outcome by either deciding to use scarce resources, such as having people work overtime or sending out a task force, towards further adapting task $s = 1$ (by choosing $m(\omega) = 1$) or towards further adapting task $s = 2$ (by choosing $m(\omega) = 2$). For simplicity, I model the manager's options in a way that task performance is perfectly adapted to the task shocks or not at all. The model's focus thereby lies on examining the manager's problem to correctly prioritise the firm's tasks and ensure that resources are funnelled to where the benefits are the greatest. This evidently abstracts from potentially interesting 'adaptation portfolios' the manager may wish to build, which may be an avenue for further research.

For the remainder of this section, I will charaterise the manager's modes of decision-making under the two feasible exposures to information: *ex ante* (, *i.e.* without regard to the current state of the economy) and *ex post* (, *i.e.* taking the current state of the economy into account). I follow the language of Dessein et al. (2016) in labelling these *management-by-plan* and *management-by-feedback*, respectively.

Management-by-plan

Definition 2. *The firm is managed by plan if the manager makes her decision solely on the basis of her fundamental knowledge such that* $\omega = \underline{\omega}$.

The manager would prefer to choose $m(\omega) = 1$ if and only if it yields higher profits than choosing $m(\omega) = 2$. Since under management-by-plan, by definition, the manager makes this decision without observing the local shocks θ_s, she must make her decision on the basis of the expected profits under each action.

$$\begin{aligned} E[-\beta_2 \cdot |\theta_2|^\lambda] &\geq E[-\beta_1 \cdot |\theta_1|^\lambda] \\ \iff \gamma &\leq 1 \end{aligned} \tag{1.3}$$

where I define $\gamma \equiv \frac{\beta_2}{\beta_1}$. I arrive at inequality 1.3 as both θ_1 and θ_2 are drawn from the same distribution.[7]

Assumption 2. *Without loss of generality, relabel β_1 and β_2 such that $\beta_1 \geq \beta_2$.*

Without observing the state-of-the economy, the manager will always choose to prioritise the task that carries the higher *ex ante* weight. By construction, this coincides with task 1.

Management-by-feedback

Definition 3. *The firm is managed by feedback if the manager observes the state of the economy before making her decision such that $\omega = \overline{\omega}$.*

When managing by feedback, the manager does not have to form expectations over the state of the economy but can base her decision where to direct the organisational focus on the severity of the actual task-level shocks. However, when incorporating this information into her decision-making she misevaluates it. Particularly, she is psychologically biased towards salient events and, as a consequence, overweights extreme outcomes that differ strongly from the expected outcomes (see Section 1.2.1 for a further discussion on this).

I closely follow Bordalo et al. (2013) in capturing this functionally. I identify the reference level for evaluating the task-level shocks with their expected value $E[\theta_s] = 0$ (see Assumption 1). Given a pair of realisations θ_1 and θ_2, the manager then inflates (deflates) the relative weight of the event that departed more (less) strongly from what was expected by a factor of $\frac{1}{\hat{\delta}}$ ($\hat{\delta}$). Hence, if she manages-by-feedback, she will choose to adapt task $s = 1$ if and only if

$$-\beta_2 \cdot \delta_2(\theta_1, \theta_2) \cdot |\theta_2|^\lambda \geq -\beta_1 \cdot \delta_1(\theta_1, \theta_2) \cdot |\theta_1|^\lambda \tag{1.4}$$

where

$$\delta_s(\theta_1, \theta_2) = \begin{cases} \hat{\delta} & \text{if } |\theta_s - E[\theta_s]| < |\theta_{-s} - E[\theta_{-s}]| \iff |\theta_s| < |\theta_{-s}| \\ 1 & \text{if } |\theta_s - E[\theta_s]| > |\theta_{-s} - E[\theta_{-s}]| \iff |\theta_s| > |\theta_{-s}| \end{cases}$$

where $0 < \hat{\delta} < 1$.

It is important to note that I assume the manager to be naive towards this bias, meaning that she is perfectly unaware of it and treats her subjective observations as

[7] Moreover, as the manager knows that whenever she does not observe any information regarding the state of the economy, neither of the task-level shocks has been classified as 'exceptional'. By construction, the adopted information restriction mechanism in this article leaves the expected value of all task-level shocks unchanged regardless of what magnitude of z the owner chooses (see Section 1.3.3 for a discussion). For this result, however, it suffices that the mechanism is the same for both θ_1 and θ_2.

truthful representations of the objective world. Moreover, the model's results are not dependant on this particular formulation of the manager's salience distortion. In Chapter B, I show that all results also go through under continuous salience distortions.

1.3.3 Critical discussion

In the current setup, either information regarding states of both tasks or neither is conferred to the manager depending on whether at least one of those states is considered 'exceptional' or not, respectively. In assessing the desirability of this form of information management, it is compared to a lucid benchmark, namely a regime of unconditional full information disclosure.

Although the owner can restrict the manager's access to information, it is important that he cannot fabricate or misrepresent it. Otherwise, the owner could anticipate the manager's bias and present the information in exactly such a way that the decision-bias and report-distortion would cancel each other out. Milgrom and Roberts (1986) have encountered the same issue in a delegation problem. I will follow them by assuming that, although information can be withheld, any information the manager does receive can be freely verified and/or there are sufficient penalties for lying.

The symmetry of the information restriction device around 0 mirrors the symmetry of the density function $g(\theta_s)$ (see Assumption 1). Hence, if in any given case the manager's information is, in fact, restricted and she must form expectations over the task-level shocks θ_s, she knows the density function's feasible support is symmetrically bounded. But even if she can perfectly observe the owner's choice of z, her unconditional expectation still remains her best bet for any possible z:

$$E[\theta_s|z] = E[\theta_s] \quad \forall z, \theta_s \quad \forall s \tag{1.5}$$

Therefore, the manager's reasoning specified in condition 1.3 still remains valid if she is aware of the fact that her information inflow is restricted in the manner described above.

The manager is unaware of her own decision-bias, meaning she acts as if she is convinced that $\hat{\delta} = 1$. Note that in this setting the manager also has no incentive to strategically threaten the owner with choosing a bad alternative in order to gain full information. As she is convinced of herself not being biased, she will expect the same judgement from the owner and interpret any attempts to restrict her information as a trembling-hand strategy, since a rational owner would always choose $z = 0$ if $\hat{\delta} \to 1$ (see Proposition 1). Second, any threats she might make in order to categorically ensure full information are not credible due to the game being one-shot and sequential. Once the owner has made his choice it is in the manager's best interest to attempt maximising profits given the information she has.[8]

Generally, the considered mechanism of flagging exception is not the only possible one. For example, the owner may consider only conveying information on the state considered exceptional or convey only information on the most exceptional state while not disclosing information regarding the other state at all. In the case of an exception actually occurring, meaning under management-by-feedback, the manager would then evaluate a certain outcome against an uncertain one when deciding whether to adapt task $s = 1$ or adapt task $s = 2$. Bordalo et al. (2012) formulate an approach determining which states attain saliency in the decision-makers mind for choices involving uncertain outcomes.

However, this approach is not suitable for the given setup for several reasons. Most importantly, it is not well-defined for decision-problems involving several dimensions, such as the two task-level shocks θ_1 and θ_2 in this paper. While in Bordalo et al. (2013), the decision-maker compares each realisation to its own reference value to determine the salient *dimension*, in Bordalo et al. (2012) the decision-makers compares realisations with each other to determine the salient *state*. It is unclear how such computations could be carried out and interpreted. Assigning saliency to states is also not specified when there are infinitely many possible states of the world. Apart from the technical inapplicability, the necessary reflections involved on the side of the decision-maker to generate saliency when (even finitely) many states exist arguably are unintuitive. Therefore, I opt to focus on full disclosure of all information in case of information transmission.

The proposed mechanism for information transmission further is intuitive and easy to implement. It can be seen as an 'alert-system', that either leaves the manager undisturbed or issues a review of the environment if potential danger is detected. In different contexts, similar mechanisms have been proposed by Reis (2004) and Andries and Had-

[8] An interesting problem arises if the owner is aware of her own decision bias, but does not know the exact value of $\hat{\delta}$. In order to improve her decision-making, she could then attempt to infer the magnitude of her bias from the owner's action and correct her perceptions accordingly. Under the assumption of the task-level shocks θ_s (see Assumption 3) being standard normally distributed, perfect knowledge of the owner's action (for all $z > 0$) is sufficient to perfectly inform her of her bias, as for given β_1, β_2 and λ (which are known by the manager), every degree of management-by-exception is uniquely rationalised by a single value of $\hat{\delta}$. If the owner does not observe z directly, she may still make inferences from information she receives given her knowledge of the structure of the decision-problem. In particular, if she does not observe any information on the state of the economy, she will infer that she must have some $\hat{\delta} < 1$, but cannot make use of this knowledge since there exist no perceptions to correct. If, on the other hand, she does receive information regarding the task-level shocks, she can infer that $z < max\{|\theta_1|, |\theta_2|\}$ and hence some lower bound on $\hat{\delta}$. Further assumptions on the manager's prior regarding her bias and an underlying distribution of it could lead to a further formalisation of this problem and may be an interesting avenue for further studies.

dad (2019). Its relative simplicity may make it particularly suitable for applied contexts while also ensuring precision in the results. These gains are traded-off with concessions in generality. In particular, considering further conditions for information disclosure may prove an interesting avenue for future research, such as non-symmetrical intervals for flagging exceptions as well as even more complex devices, but are considered outside the scope of this paper.

1.3.4 Construction of indirect profit function

Figure 1.2 depicts the manager's choice correspondence under all possible scenarios of the world. Anticipating this, the owner can express the expected profits as a tractable function of z. Due to the symmetry with respect to θ_1 and θ_2 of both the profit and the density function $g(\cdot)$ around 0, the bounds of the integrals can be mirrored arbitrarily around both θ_s axes. We can thus restrict the focus of the investigation to positive values of θ_1 and θ_2. I arrive at the following expression:

$$\begin{aligned} E[\pi(z)] = 4 \cdot & \int_0^\infty \int_0^{max\{\left(\frac{\hat{\delta}}{\gamma}\right)^{\frac{1}{\lambda}} \cdot \theta_1, \theta_1\}} g(\theta_1) \cdot g(\theta_2) \cdot (-\beta_2 \theta_2^\lambda)\, \mathrm{d}\theta_2\, \mathrm{d}\theta_1 \\ +4 \cdot & \int_0^\infty \int_{max\{\left(\frac{\hat{\delta}}{\gamma}\right)^{\frac{1}{\lambda}} \cdot \theta_1, \theta_1\}}^\infty g(\theta_1) \cdot g(\theta_2) \cdot (-\beta_1 \theta_1^\lambda)\, \mathrm{d}\theta_2\, \mathrm{d}\theta_1 \\ +4 \cdot \int_0^{min\{\left(\frac{\gamma}{\hat{\delta}}\right)^{\frac{1}{\lambda}} \cdot z, z\}} & \int_{max\{\left(\frac{\hat{\delta}}{\gamma}\right)^{\frac{1}{\lambda}} \cdot \theta_1, \theta_1\}}^z g(\theta_1) \cdot g(\theta_2) \cdot (\beta_1 \theta_1^\lambda - \beta_2 \theta_2^\lambda)\, \mathrm{d}\theta_2\, \mathrm{d}\theta_1 \end{aligned}$$

Figure 1.2: Decision tree of the model

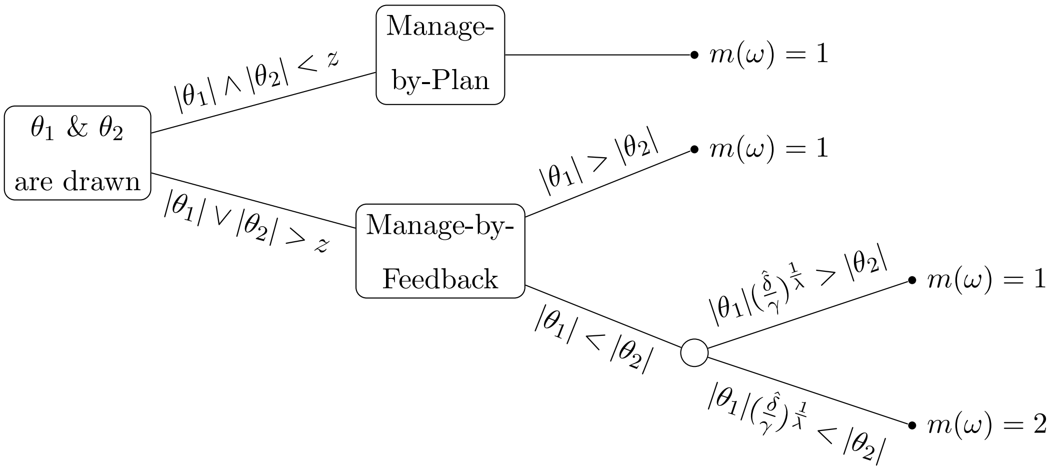

Management-by-exception implemented by flow of information. Under "normal" conditions, the manager adheres to the *ex ante* optimal plan without regard to the state of the economy. If sufficiently unusual events occur, a briefing is called and the firm is managed by feedback.

As Figure 1.2 shows, the manager only chooses to adapt task $s = 2$ if she manages-by-feedback and both $|\theta_1| < |\theta_2|$ and $|\theta_1|(\frac{\hat{\delta}}{\gamma})^{\frac{1}{\lambda}} < |\theta_2|$ hold. Of the latter two conditions, only one can be binding at any given time. For the sake of analytical simplicity, I can therefore replace the $min\{\cdot\}$- and $max\{\cdot\}$-operator by substituting the $\hat{\delta}$ in the expression above with the following parameter:

Definition 4. $\delta := max\{\gamma, \hat{\delta}\}$

Importantly, values for $\hat{\delta}$ lower than γ are not excluded from the analysis, but simply do not have any effect on it. I can then take the derivative of $E[\pi(z)]$ with respect to z. Applying Leibniz' Rule of Integration twice yields the following expression of the slope of the expected profits with respect to the implemented degree of *management-by-exception*:

$$\frac{\partial E[\pi(z)]}{\partial z} = 4 \cdot g(z) \cdot \int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1) \cdot (\beta_1 \theta_1^\lambda - \beta_2 z^\lambda) \, \mathrm{d}\theta_1 \tag{1.6}$$

1.4 Management-by-Exception as a Dominant Strategy

In this section, I theoretically map out relevant factors in the adoption choice of the management practice *management-by-exception*, yielding testable predictions while also potentially shedding light on earlier findings. To this end, I conduct the investigation with an inherently binary view on the owner's action space, by determining sufficient conditions when he will optimally choose a strictly positive value for z as opposed to choosing $z = 0$. By Definition 1 this corresponds to *management-by-exception* being implemented at least to some degree. Although this analysis lacks precision, it retains generality, as it is done under the minimal assumptions specified in Assumption 1. In the following section I will show that maximum precision can be obtained with further assumptions on the distribution of the economy's state.

First, it is helpful to establish some properties of the following function:

$$k(z) \equiv \frac{z \cdot g((\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z)}{\int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1)\, \mathrm{d}\theta_1} \tag{1.7}$$

defined on the domain $z \in (0, \infty)$.

Lemma 1. *The function k*

(i) is continuous.

(ii) satisfies $\lim_{z \to 0} k(z) = \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \geq 1$.

Before proceeding, note that since the density function $g(\,\cdot\,)$ is, by Assumption 1, continuous and symmetric around 0, it holds that $4 \cdot g(z) > 0$ in the vicinity of $z = 0$. Therefore $\frac{\partial \pi}{\partial z} \geq 0$ if and only if

$$f(z) := \int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1) \cdot (\beta_1 y_1^{\lambda} - \beta_2 z^{\lambda})\, \mathrm{d}\theta_1 \geq 0 \tag{1.8}$$

Characterising $f(z)$ is thus equivalent to characterising the gradient of the expected profits in the vicinity of $z = 0$. The following finding provides a tool for doing so.

Lemma 2. *For $z > 0$ and $1 > \gamma > 0$, $f'(z) \geq 0$ if and only if* $k(z) \geq \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$.

Before deriving the main result, it is helpful to make two observations. First, by an application of the Combination Theorem for Continuous Functions, the slope of the profit function is continuous at all $z \in [0, \infty)$. Hence, although the function $k(z)$ is not defined at $z = 0$ itself, the function of main interest $\frac{\partial E[\pi(z)]}{\partial z}$ on which inferences are made is defined and well-behaved at this point. Second, note that $\frac{\partial E[\pi(z)]}{\partial z}|_{z=0} = 0$ for all feasible settings. Therefore, if the slope of the expected profits increases when increasing z from 0, it becomes positive. In this case, expected profits also rise by increasing z from $z = 0$, rendering $z = 0$ (not implementing *management-by-exception*) a dominated alternative. If, on the other hand, the slope of the expected profits decreases at all strictly positive z, the owner would prefer choosing $z = 0$. Using these insights, it is easy to first verify a rather technical insight.

Proposition 1. *Management-by-exception is only adopted if there is a behavioural bias in the manager's decision-making ($\widehat{\delta} < 1$).*

This stems from the fact that, within the model, the manager's inclination to misevaluate her available information is the sole driver for implementing *management-by-exception*. Specifically, her behavioural bias causes her to misprioritise the firm's tasks and fail to direct the organisational focus to the most important areas. By effectively introducing rigid plans on *ex ante* optimal strategies, the manager's misprioritisation through 'over-responsiveness' can be reversed, however, at the cost of misprioritisation through 'under-responsiveness'. If for some ranges of moderate realisations of the task-level shocks θ_1 and θ_2 the first effect dominates the second, *management-by-exception* at least to some degree improves prioritisation and should thus be adopted.

As stated earlier, the model does not claim exclusiveness. There may well be further motivations for implementing the management technique, such as costly information transmission. It appears to be a relevant objective, though, since the practician (Mackintosh, 1978, p. 96) in the chapter *"Activating the system that spotlights the biggest problems"* concerning *management-by-exception* states: *'Many times management will commit the entire company to a drive for the elimination of what it conceives to be major operating problems. All too often, however, these drives are initiated with inadequate or partial information (...), and result in little more than a futile, misguided expenditure of time and effort.'*

Crucially, using the insights from Lemma 1 and Lemma 2, it is possible to derive a condition for when *management-by-exception* should be adopted.

Proposition 2. *For all distributions of the task-level shocks θ_1 and θ_2 satisfying Assumption 1 the firm adopts management-by-exception if $\lambda \cdot \frac{\delta}{1-\delta} < 1$.*

From this result, I will now state and discuss the comparative statics with respect to the model's underlying parameters. For all following Corollaries, I shall state that an adoption of *management-by-exception* becomes 'more likely' if a change in a certain parameter will cause the condition specified in Proposition 2 to hold for a strictly larger range of the remaining parameters.

Corollary 1. *Ceteris paribus, a stronger bias in the manager's decision making (lower $\widehat{\delta}$) makes an adoption of management-by-exception more likely.*

This result provides a possible explanation for some previous, puzzling empirical findings. It has long been surmised that *'in high-risk conditions where safety is of concern, active management-by-exception may play a more prominent and effective role'* (Antonakis et al., 2003, p. 270). Bass and Avolio (2000) corroborate this by studying the performance of U.S. Army platoons. They find large positive effects of *management-by-exception* on a unit's readiness. This may be a result of leaders finding it difficult to maintain cold and objective decision-making in contexts that are fast-paced and potentially life-threatening. The more this is the case, the more they may benefit from rigid protocols that are followed irrespective of the situation's developments unless sufficiently exceptional circumstances arise.

Apart from contextual factors that may lead to the adoption of *management-by-exception*, there may be manager-specific determinants, too. (Antonakis et al., 2003, p. 274) remark that *'men tend to use management-by-exception more often than do women'*. It is a well-established finding that men are generally more overconfident than women (e.g. Barber and Odean, 2008). The stronger prevalence of the management technique may be due to mirroring, gender-specific phenomena regarding overreactions to salient stimuli.

Corollary 2. *Ceteris paribus, a stronger valuation of extreme events in the manager's decision making (higher λ) makes an adoption of management-by-exception less likely.*

Perhaps surprisingly, although *management-by-exception* appears to do well in dangerous, high-stakes environments, a higher weight on extreme outcomes makes the technique less attractive for the firm. If outliers become more important, then a more sensitive policy (,*i.e.* a lower z) will be desired so that more events can be dealt with individually rather than via a rigid plan. Although some studies such as Bernile et al. (2017) tie corporate policies back to risk attitudes, surprisingly little research has been done examining the adoption choice of *management-by-exception* towards such sentiments.

Corollary 3. *Ceteris paribus, a stronger polarisation in the ex ante importance across the firm's tasks (lower γ) makes an adoption of management-by-exception more likely.*

The desirability of *management-by-exception* increases if the *ex ante* quality of the plan rises. To illustrate this point, consider a firm where $\beta_2 \to 0$ so that $\gamma \to 0$. It will then never be optimal for the manager to adapt task $s = 2$. *Management-by-plan* will then always select the best strategy, thus increasing the desirability of a highly rigid *management-by-exception* policy. This result reflects findings from Dessein et al. (2016); in a similar functional setting where there are limits on information transmission within the firm, the authors find that an increased *ex-ante* importance of already important tasks may make it desirable to make the organisation more rigid by managing more of the firm's tasks by plan. However, while in their setup this outcome is static and independent of the economy's state, under *management-by-exception* the firm will always revert to *management-by-feedback* if sufficiently exceptional circumstances arise.

Since with the *Multifactor Leadership Questionnaire* survey there already exists a highly used, common tool to elicit the model's main variable, the major challenge for empirical estimations of the model's predictions lies in identifying and constructing measures of the key parameters γ, $\hat{\delta}$ and λ. Firm's accounting posts of different divisions may provide proxies for imbalances of costs across its products' dimensions (γ). Concerning the remaining drivers, there is a large literature in linking the adoption of management practices to psychometric indicators of the leaders (see Bono and Judge, 2004, for a meta-analysis). Moreover, in 2015 the updated version of the large-scale *Management and Organizational Practices Survey* has acknowledged the importance of proper decision-making as well as utilising correct information inputs to do so, by including the new section *data and decision-making.* In a similar vein, eliciting personal or contextual risk- and salience sentiments may prove insightful.

1.5 Existence and Uniqueness of Optimal Degree of Management-by-Exception

In this section, I will show that what the firm should treat as '*exceptional*' and what as '*usual*' can be precisely defined if one is willing to place further assumptions on the distribution function of the state of the economy $g(\,\cdot\,)$.[9] For the remainder of the article

[9]The conception and formulation of this proof has greatly benefited from extensive discussions with Bram Driesen.

I assume that the task-level shocks θ_1 and θ_2 are independently drawn from a standard normal distribution:

Assumption 3. *Let*

$$g(\theta_s) \equiv \frac{1}{\sqrt{2\pi}} e^{\frac{-\theta_s^2}{2}} \equiv \phi(\theta_s) \quad \forall s \tag{1.9}$$

As this assumption constitutes a special case of the previous set placed on the distribution of the task-level shocks (Assumption 1), all results from Proposition 2 still obtain under the current setting. However, it is now possible to derive additional properties of the function $k(z)$ (see Equation (1.7)).

Lemma 3. *The function k*

(i) is continuous.

(ii) satisfies $\lim_{z\to 0} k(z) = \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \geq 1$.

(iii) satisfies $\lim_{z\to\infty} k(z) = 0$.

(iv) is strictly decreasing in z.

The first two items ensure that under certain conditions there is a segment where the firm's profit rises when z is increased marginally from 0, yielding the previous section's main result. Items three and four further ensure that either the firm's profits monotonically decrease with z, or the profit function can be separated into a rising segment (for low z) and a decreasing segment (for high z). These findings are also illustrated by Figure A.1 in the appendix.

Given these insights, it is possible to obtain the following result.

Proposition 3. *If the task-level shocks θ_1 and θ_2 are standard normally distributed, there exists a uniquely optimal degree of implementing management-by-exception for all possible firms.*

In order to develop intuition for the qualitative and quantitative behaviour of the model with respect to its underlying parameters, I conduct a short simulation exercise. Although a closed-form solution for z does not exist, the owner's optimal action can be solved for numerically with arbitrary precision. Due to the additional insight in Lemma 3 *(iv)*, Lemma 2 now dissects the owner's expected profit function with respect to z into either a single falling and a single rising segment or instead causes it to be falling at all $z > 0$. In any case, the function will be single-peaked. Therefore, the optimal z must

Figure 1.3: Simulated illustration of optimal organisational responses with respect to the model's deep parameters.

(a) $\hat{\delta} = 0.25$

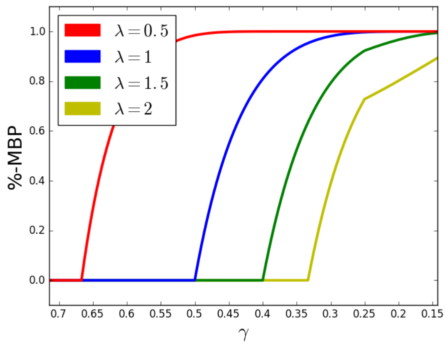

(b) $\hat{\delta} = 0.35$

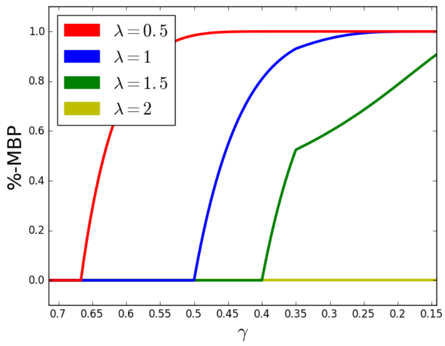

(c) $\hat{\delta} = 0.45$

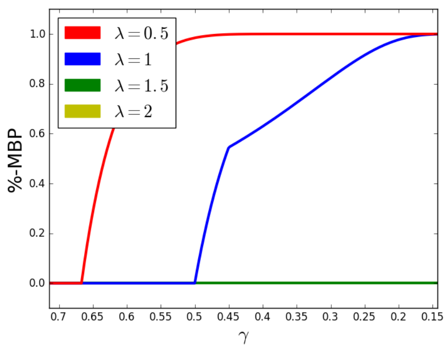

(d) $\hat{\delta} = 0.55$

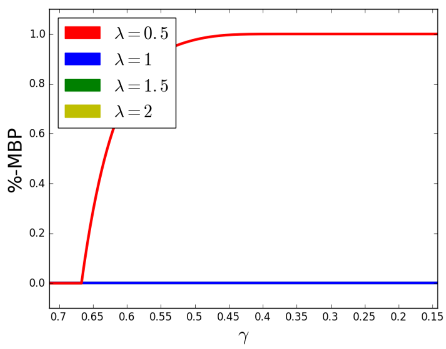

Increasing firm rigidity, measured by the optimal share of activities managed by plan (see Equation (1.10)), with increasing overresponsiveness of the manager ($\hat{\delta}$), stronger polarisation of *ex ante* importance of tasks (γ) and stronger convexity of costs with respect to shocks (λ).

be located between two values where the slope of the expected profit function changes its sign. Starting from 0, I simulate this slope for all z in small steps and pick the first z where a change of sign occurs as the optimal solution.

To further facilitate interpretation, I express the owner's optimal solution as the expected share of the firm's activities dealt with via *management-by-plan*, henceforth denoted $\%MBP(z^*)$:

$$\begin{aligned} \%MBP(z^*) &\equiv \Phi^{2-sided}_{\theta_1,\theta_2}(z^*) = \Phi^{2-sided}_{\theta_1}(z^*) \cdot \Phi^{2-sided}_{\theta_2}(z^*) \\ &= \int_{-z^*}^{z^*} \phi(\theta_1)\, \mathrm{d}\theta_1 \cdot \int_{-z^*}^{z^*} \phi(\theta_2)\, \mathrm{d}\theta_2 \end{aligned} \tag{1.10}$$

This formulation intuitively picks up on the notion that higher levels of z correspond to a more rigid organisation (see Definition 1). As z increases, the share of events dealt with via a static plan without regard to the state of the world also increases, whereas at

$z = 0$ all management is done case-by-case.

Figure 1.3 illustrates the solutions of Equation (1.10) for some exemplary parametric settings. Echoing the findings from Proposition 2, for each λ there exists a double threshold with respect to $\hat{\delta}$ and γ above which *management-by-exception* will not be implemented. However, as soon this threshold is crossed, the firm's drive to implement the practice rises sharply. For example, for $\lambda = 1$ and the manager's bias with $\hat{\delta} = 0.45$ just below the required threshold the firm already executes around 80% of its activities by plan if one task is three times more important than the other, making the firm very rigid. Although the exact values elicited should be taken with care, they suggest that the investigated underlying factors may, in fact, shape organisational structures in an economically meaningful way.

1.6 Conclusion

Although *management-by-exception* is a well-established concept enjoying widespread use in practice in a large variety of fields, this paper is the first to lay out a mathematical framework in order study the technique with theoretical rigour. I derive two main conclusions. First, under relatively mild assumptions I obtain a simple, sufficient condition for when firms will adopt *management-by-exception*. Second, I show that there exists a uniquely optimal degree of implementing *management-by-exception* with respect to the underlying parameters if the state of the economy is standard normally distributed.

Testing the model's predictions is an obvious next step for future research (see Section 1.4 for a discussion). Theoretically, it may be interesting to see how a growing number of tasks, interpretable as a growing firm size, would affect the adoption choice of *management-by-exception* as well as how these drivers would interact with the ones established in this paper. Further, it may be interesting whether other established psychological biases, such as overconfidence or confirmatory bias, on the part of the manager could act as drivers in adopting *management-by-exception*.

Appendix A

Proofs

A.1 Proof of Lemma 1

(*i*) As all of z, $g(z)$ and $\int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1)\,\mathrm{d}\theta_1$ are continuous functions everywhere, by the Combination Theorem for Continuous Functions it holds that $k(z)$ is also continuous everywhere.

(*ii*) As $k(0) = \frac{0}{0}$ and both its numerator and denominator are differentiable everywhere, l'Hôspital's rule may be applied.

$$\lim_{z\to 0} k(z) = \lim_{z\to 0} \frac{\frac{\partial}{\partial z}\left[z \cdot g((\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z)\right]}{\frac{\partial}{\partial z}\left[\int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1)\,\mathrm{d}\theta_1\right]} = \lim_{z\to 0} (\frac{\delta}{\gamma})^{\frac{1}{\lambda}} + z \cdot \frac{g'\left((\frac{\gamma}{\delta})^{\frac{1}{\lambda}} \cdot z\right)}{g\left((\frac{\gamma}{\delta})^{\frac{1}{\lambda}} \cdot z\right)} = \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \tag{A.1}$$

Because of Definition 4 it must hold that $\left(\frac{\delta}{\gamma}\right) \geq 1$. Since $\lambda > 0$, we must have that $\lim_{z\to 0} k(z) \geq 1$.

A.2 Proof of Lemma 2

Differentiating $f(z)$ with respect to z yields

$$f'(z) = \left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} \cdot g\left(\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} \cdot z\right) \cdot z^{\lambda}\left(\beta_1\left(\frac{\gamma}{\delta}\right) - \beta_2\right) - \beta_2 \lambda z^{\lambda-1} \cdot \int_0^{(\frac{\gamma}{\delta})^{\frac{1}{\lambda}} z} g(\theta_1)\,\mathrm{d}\theta_1 \tag{A.2}$$

Then for $z > 0$ and $1 > \gamma > 0$, setting $f'(z) \geq 0$ and rearranging yields the desired condition.

A.3 Proof of Proposition 1

Since, in this paper, the manager's behavioural bias is the sole driver for the owner's motivation to adopt *management-by-exception*, doing so can never be desirable when no such bias exists. It is easy to verify this. By Lemma 2, $f(z)$ and thus $\frac{\partial E[\pi]}{\partial z}$ decreases if and only if $k(z) < \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$. If the manager's decision bias vanishes as $\widehat{\delta} \to 1$ (and thus $\delta \to 1$), the right-hand side goes to ∞ and the inequality must hold for all z. Coupled with the insight that $\frac{\partial E[\pi]}{\partial z}|_{z=0} = 0$, this implies that in this case $\frac{\partial E[\pi]}{\partial z} < 0$ for all z and that the owner can increase expected profits by decreasing z from all strictly positive levels.

A.4 Proof of Proposition 2

Recall that by Lemma 1 *(ii)* $\lim_{z\to 0} k(z) = \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}}$. When $z \to 0$ it then holds that $k(z) > \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ if and only if $1 > \lambda \cdot \frac{\delta}{1-\delta}$. Due to the parsimonious assumptions placed on $g(\,\cdot\,)$ so far, not much can be said about the behaviour of $k(z)$ as z increases. Assume, however, that at $\tilde{z}$ it holds that $k(\tilde{z}) = \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$, where $\tilde{z}$ may be arbitrarily large. If then $1 > \lambda \cdot \frac{\delta}{1-\delta}$ holds with strict inequality and since by Lemma 1 *(i)* $k(z)$ is continuous, there must be a range $(0, \tilde{z})$ with some strictly positive length where it also holds that $k(z) > \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$. By Lemma 2, this also implies that $f'(z) > 0$ for the entire range. Because $\frac{\partial E[\pi]}{\partial z}|_{z=0} = 0$, this further implies that $\frac{\partial E[\pi]}{\partial z} > 0$ for all $z \in (0, \tilde{z})$. In this case, starting from $z = 0$ the owner can then increase expected profits by increasing z up to $\tilde{z}$, rendering $z = 0$ a strictly dominated strategy for the owner. If, on the other hand, there should exist no $\tilde{z}$ such that $k(\tilde{z}) = \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ (still assuming that $1 > \lambda \cdot \frac{\delta}{1-\delta}$), then $k(z) > \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ for all z and all $z > 0$ are better than $z = 0$. Therefore, if $1 > \lambda \cdot \frac{\delta}{1-\delta}$ then not implementing management-by-exception to any degree by choosing $z = 0$ is a strictly dominated strategy.

A.5 Proof of Corollaries 1-3

The comparative statics results can be directly gained from the dominance condition in Proposition 2. Not implementing *management-by-exception* is a strictly dominated option if $\lambda \cdot \frac{\delta}{1-\delta} < 1$. *Ceteris paribus*, this condition is more likely to hold if either λ or δ decrease. Recall, however, that by Definition 4, δ is bounded from below by the maximum of the parameters $\widehat{\delta}$ and γ. Hence, for a decrease in one of the two parameters to affect the adoption choice, the other one must be sufficiently low.

A.6 Proof of Lemma 3

(i) See Proof of Lemma 1.

(ii) See Proof of Lemma 1.

(iii) To show the desired result, I determine the limits of the numerator and denominator of $k(z)$ individually. For the denominator, note that because of symmetry of $\phi(\,\cdot\,)$ coupled with the fact that $\phi(\,\cdot\,)$ is a density function, it must hold that $\int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}}z} \phi(\theta_i)\,\mathrm{d}\theta_i = \frac{1}{2} > 0$.

For the numerator, first note that it must hold that $\lim_{z\to\infty} \int_0^z \phi(\theta_i)\theta_i\,\mathrm{d}\theta_i < \infty$, since by symmetry of $\phi(\,\cdot\,)$ (Assumption 1 *(iii)*) we would otherwise also have that $\lim_{z\to-\infty} \int_z^0 \phi(\theta_i)\theta_i \mathrm{d}\theta_i = -\infty$. In that case, however, $E[\theta_i]$ would not exist, contradicting Assumption 1 *(i)*.

With this in mind, I will prove the desired result by contradiction, assuming that $\lim_{z\to\infty} k(z) = C > 0$ where C is some finite, strictly positive constant. Then for all $\varepsilon > 0$, there exists a positive number M such that for all $\theta_i > M$ it holds that $|\phi(\theta_i)\cdot\theta_i - C| < \varepsilon$. Therefore, for all $\theta_i > M$ it must also hold that $C-\varepsilon < \phi(\theta_i)\cdot\theta_i < C+\varepsilon$. However, then for all $\varepsilon < C$ we would also have that

$$\infty = \int_M^\infty C - \varepsilon\,\mathrm{d}\theta_i < \int_M^\infty \phi(\theta_i)\theta_i\,\mathrm{d}\theta_i < \int_0^\infty \phi(\theta_i)\theta_i\,\mathrm{d}\theta_i$$

which contradicts the initial finding. As C cannot be negative, I conclude that it must hold that $C = 0$.

Combining the two insights, I conclude that

$$\lim_{z\to\infty} k(z) = \frac{0}{\frac{1}{2}} = 0 \tag{A.3}$$

(iv) Define the function

$$l(z) := \left(\int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}}z} \phi(\theta_1)\,\mathrm{d}\theta_1\right) \cdot \left(1 - \left(\frac{\gamma}{\delta}\right)^{\frac{2}{\lambda}} \cdot z^2\right) - z \cdot \phi\left(\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z\right) \cdot \left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} \tag{A.4}$$

Note that $l(0) = 0$ and $l'(z) = -2z \cdot \left(\frac{\gamma}{\delta}\right)^{\frac{2}{\lambda}} \cdot \left(\int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}}z} \phi(\theta_1)\,\mathrm{d}\theta_1\right) < 0$ for all $z > 0$. Hence, it must hold that $l(z) < 0$ for all $z > 0$.

Figure A.1: Illustration of findings from Lemma 2 and Lemma 3.

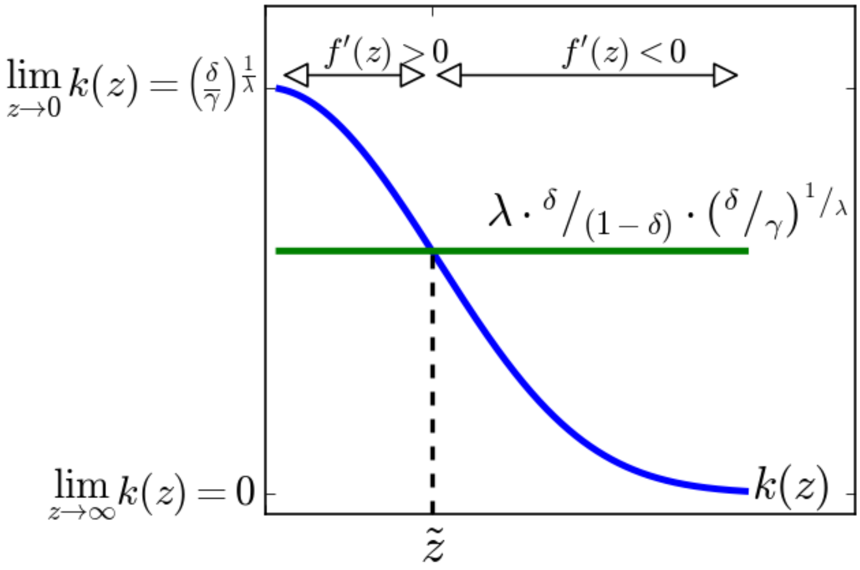

Differentiating $k(z)$ with respect to z yields:

$$k'(z) = \frac{l(z)\phi\left(\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z\right)}{\left(\int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z} g(\theta_1)\, \mathrm{d}\theta_1\right)^2} \tag{A.5}$$

Because $\phi\left(\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z\right) > 0$, $\left(\int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z} g(\theta_1)\, \mathrm{d}\theta_1\right)^2 > 0$ and $l(z) < 0$ for all $z > 0$, it holds that $k'(z) < 0$ for all $z > 0$.

A.7 Proof of Proposition 3

The proof builds on the findings from Lemma 2 and Lemma 3, which are illustrated in Figure A.1. It follows an intermediate-value-theorem-type of argument, relating $\lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ and $k(z)$ (see Equation (1.7)). Since the former is a constant and the latter, by Lemma 3 *(i)* and *(iv)*, is continuous and strictly decreasing in z, there can either be no crossing or exactly one crossing between the two. I will discuss both possible cases in order.

As in Proposition 2, when $z \to 0$, by Lemma 3 *(ii)*, $k(z) < \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ if and only if $\lambda \cdot \frac{\delta}{1-\delta} > 1$. However, since by Lemma 3 *(iv)* $k(z)$ is strictly decreasing in z when θ_1 and θ_2 are drawn from a standard normal distribution, then this also implies that $k(z) < \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ for all z. In turn, by Lemma 2 this implies that $f'(z) < 0$ for all z, which, coupled with the fact that $f(0) = 0$, implies that $f(z) < 0$ for all z. As characterising the sign of $f(z)$ is equivalent to characterising the slope of the expected profits, we know from this that $\frac{\partial E[\pi]}{\partial z} < 0$ for all z. Therefore, if $\lambda \cdot \frac{\delta}{1-\delta} > 1$ the owner can increase expected profits by decreasing z from all strictly positive levels. The uniquely optimal level for z is then at the its lowest possible level $z = 0$.

If, on the other hand, $\lambda \cdot \frac{\delta}{1-\delta} < 1$, then it holds that $\lim_{z\to 0} k(z) > \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta} > \lim_{z\to\infty} k(z)$, where the first inequality follows from Lemma 3 *(ii)* and the second from Lemma 3 *(iii)* (for all feasible parametric settings). Continuity of $k(z)$ (Lemma 3 *(i)*), by the intermediate value theorem, ensures that there is at least one crossing between $k(z)$ and $\lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ while the fact that $k(z)$ is strictly decreasing $k(z)$ (Lemma 3 *(iv)*) ensures that there is not more than one crossing. Denote $\tilde{z}$ where $k(\tilde{z}) = \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$. Thus, for all $z \in [0, \tilde{z})$ it holds that $k(z) > \lambda \cdot \left(\frac{\delta}{\gamma}\right)^{\frac{1}{\lambda}} \frac{\delta}{1-\delta}$ and by Lemma 2 that $f'(z) > 0$. Since furthermore $f(0) = 0$, it then holds that $f(z) > 0$ and thus that $\frac{\partial E[\pi]}{\partial z} > 0$ for all $z \in [0, \tilde{z})$. As in Proposition 2, the owner can strictly increase the firm's expected profits by increasing z at all $z \in [0, \tilde{z})$ whenever $\lambda \cdot \frac{\delta}{1-\delta} < 1$. However, because of the single-crossing property, it now also holds that the slope of the expected profits decreases for all $z \in (\tilde{z}, \infty)$.

In order to ensure that a maximum is reached for some z, we now simply need to ensure that $f(z)$, and therefore the slope of the expected profits, does indeed turn negative at some point as z increases. This can be verified as follows.

$$\lim_{z\to\infty} f(z) = \beta_1 \lim_{z\to\infty} \int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z} \phi(\theta_1)\theta_1^{\lambda}\, \mathrm{d}\theta_1 - \beta_2 \lim_{z\to\infty} z^{\lambda} \int_0^{\left(\frac{\gamma}{\delta}\right)^{\frac{1}{\lambda}} z} \phi(\theta_1)\, \mathrm{d}\theta_1 = -\infty \quad \text{(A.6)}$$

As noted in the proof for Lemma 3 *(iii)*, $\lim_{z\to\infty} \int_0^z \phi(\theta_1)\theta_1\, \mathrm{d}\theta_1$ converges to a finite constant. Since $0 < \delta, \gamma, \lambda < \infty$, so must the first addend. It is then easy to see that the second addend goes to $-\infty$, showing the desired result. Therefore, there must be some $z^* > \tilde{z}$ for which it holds that $\frac{\partial E[\pi]}{\partial z} > 0$ for all $z \in [0, z^*)$ and $\frac{\partial E[\pi]}{\partial z} < 0$ for all $z \in (z^*, \infty)$. Therefore, there also exists a unique level of z that maximises the firm's expected profits when $\lambda \cdot \frac{\delta}{1-\delta} < 1$, proving the desired result.

Appendix B

Robustness to Continuous Salience Distortions

For simplicity, I work with rank-based salience distortions in this paper, so that the manager's evaluation becomes deflated by a constant proportion for the less salient option. However, the results are not dependant on this exclusive formulation. Alternatively, the manager's decision distortion may also be stronger for increasingly unusual shocks she perceives.

Consider a manager that manages by feedback, choosing to adapt task 1 (m=1) if and only if

$$-\beta_2 \cdot |\theta_2|^{(1-\delta)} \cdot |\theta_2|^{\lambda} \geq -\beta_1 \cdot |\theta_1|^{(1-\delta)} \cdot |\theta_1|^{\lambda} \tag{B.1}$$

This constitutes an alternative formulation to condition (1.4) in the main paper, inspired by Appendix C of Bordalo et al. (2013).

All of the paper's Lemmas and Propositions still hold in this case. When constructing the indirect profit function, no redefinition of the distortion parameter δ is necessary, as now there are no case distinctions in determining which option receives the distortion. A segment where $\frac{\partial E[\pi(z)]}{\partial z} > 0$ is then introduced whenever

$$\frac{z \cdot g(\gamma^{\frac{1}{\lambda+(1-\delta)}} z)}{\int_0^{\gamma^{\frac{1}{\lambda+(1-\delta)}} z} g(\theta_1)\, \mathrm{d}\theta_1} \geq \lambda \cdot \left(\frac{1}{\gamma}\right)^{\frac{1}{\lambda+(1-\delta)}} \cdot \frac{\gamma}{\gamma^{\frac{\lambda}{\lambda+(1-\delta)}} - \gamma} \tag{B.2}$$

It is easy to verify that the condition's right-hand side goes to infinity both when $\delta \to 1$ and when $\gamma \to 1$.

Chapter 2

Information Management against Excessive Stock Trading: More or Less? Or Both?

2.1 Introduction

Private households that participate in the stock market trade a lot. Excessive trading, through its various associated costs such as commission fees and bid-ask spreads, incurs sizeable losses. Barber et al. (2008) estimate that in Taiwan such losses are equivalent to 2.2% of gross domestic product. Similar behaviour has been documented in the U.S. (Barber and Odean, 2000), Canada (Linnainmaa et al., 2018) and Sweden (Dahlquist et al., 2016).[1] Despite this being a prevalent and well-documented phenomenon, little research has been done on how to counter this behaviour. Some scholars (see Huber et al., 2012; Hanke et al., 2010) have explored the impact of 'Tobin-like' taxes on market outcomes and individual behaviour in various settings. To the best of my knowledge, this article is the first to examine information management as a tool to curb excessive trading.

In particular, I present evidence from a computer laboratory experiment. In the experiment, participants engage in speculative trading tasks on the basis of real-life stock data where they are incentivised to realise the highest possible returns. In each instance, they can either choose to keep an allocated stock and receive its returns or pay a commission fee in order to swap their stock for another one if they fear that the

[1] The authors examine individual investors that trading the umbrella of a Swedish Pensions Agency, where all transaction costs are collectively borne by the fund. Although highly active investors individually outperform inactive ones, the overall impact of such active trading on the fund's net returns is negative and large.

stock will incur large losses. To inform this choice, participants may receive the return performance history of their allocated stock.

Figure 2.1 shows that in this context the same qualitative phenomenon noted in earlier work emerges: when ordering investors into quintiles according to trading activity, no clear pattern emerges regarding the gross returns earned. However there is a clear monotonic trend of decreasing net returns (gross returns minus trading costs) with increasing trading activity. The aim of this paper is to explore the propensity of two pure information interventions in curbing excessive trading and increasing net returns.

Figure 2.1: Stylized replication of Barber and Odean (2000).

(a) Evidence from Barber and Odean (2000)

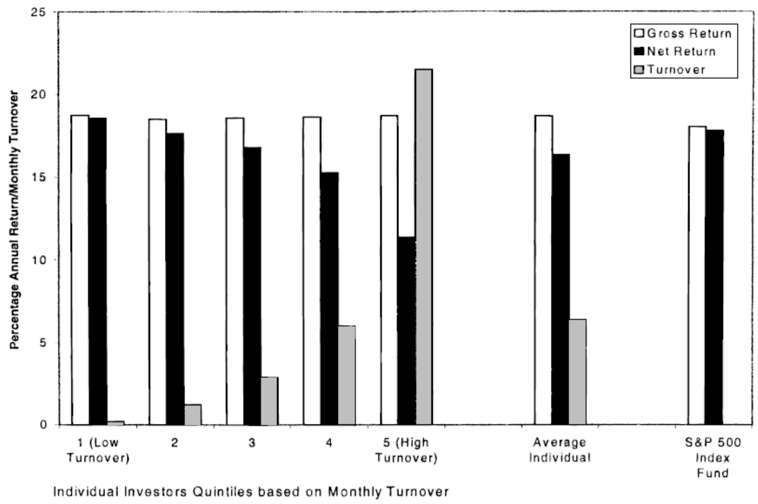

(b) Experimental evidence.

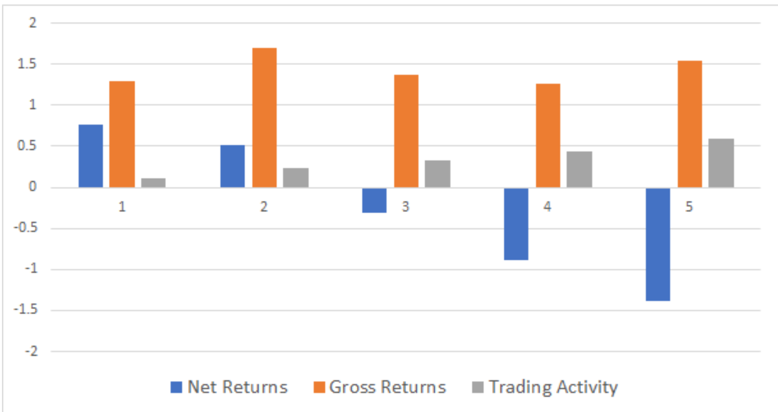

Several studies during the last two decades have made progress in explaining the 'excessive trading puzzle' both theoretically and empirically. One strand of literature attempts to explain the behavioural pattern as an outcome of a rational decision-making process. Seru et al. (2009) argue that costs associated with high trading volumes are a symptom explorative learning by unexperienced traders. On the hand, this includes familiarisation with a new environment, such as establishing thorough knowledge of all relevant parameters and calibrating prediction models with increased experience, in order to improve the expected returns from trading over time. On the other hand, this includes familiarisation with one's own ability. Grinblatt et al. (2012) show that the returns of individual traders vary greatly with their inherent ability, measured by their IQ. For excessive trading in particular, Cronqvist and Siegel (2014) find from a twin study that genetics account for a large share of variation in net returns from trading when accounting for other factors. Traders who discover their own ineptitude from past experience may then *"stop trading actively, choosing instead to invest in a passive investment"* (Seru et al., 2009, p. 706). While the authors indeed find that this second type of learning is the dominant driver of return improvements over time in their sample, Barber et al. (2020) contest this view by showing that losses from trading are both prevalent and persistent in a separate dataset.

I introduce a treatment to test this hypothesis, henceforth called the Detailed Feedback Treatment. In previous observational studies, traders voluntarily opt into gathering experience by deciding to trade more or less actively. While this choice itself may be driven by the trader's underlying ability, it may also be driven by a host of other factors. This complicates a clean estimation of the effect of experience on trading outcomes, in particular trading activity and returns. In a controlled experimental environment, it is possible to randomly administer experience by providing participants with detailed feedback on how their actions impact the outcomes of interest at every stage, what their outcome would have been if they had chosen differently as well as a prompt of whether they have chosen the most profitable option in a given instance. This enrichment of the participants' information sets increases the salience of key parameters in the decision-making environment and makes it easy to discover prolonged poor performances caused by overly active trading.

Another strand of literature regards excessive trading as a mistake driven by psychological distortions in perception. Odean (1998a) presents a model where investors overweight incoming information and underweight their prior information. This additional information may then cause overreactions to changes in the environment of the investors and consequently lead to inflated trading volumes. Previous studies have compiled evidence corroborating this view both in the laboratory and with real-life transactional data (see section 2.4.1 for a further discussion). It remains elusive, however, how such investors could improve their returns from trading. Increased experience, as suggested by Seru et al. (2009), may not lead to the acquisition of correct knowledge if the learning models of the traders are misspecified.

I propose a novel, second treatment to address this question, henceforth called the Information Filter Treatment. Instead of providing the participants with further information, I reduce the participants' access to information on changes in their environment, namely the information of their allocated stock's previous performance. Although common wisdom dictates that, in individual decision-making, more information cannot lead to worse outcomes, this need not hold true for individuals who make mistakes in processing information. In this scenario, incoming information may 'tempt' investors away from the correct. For example, while in many cases a simple buy-and-hold strategy may be most profitable for amateur investors, individuals may fall victim to panic-sales if they observe a recent history of bad returns or try to capitalise on upswings. These investors may then improve their trading outcomes, by being shielded from such information.

Both experimental treatments individually can provide insight on the potential causes of overtrading and how to counter them effectively. It remains unclear, however, whether the two respective (potential) causes for excessive trading are independent channels or to which extent they overlap. On one hand, individuals incorrectly evaluating incoming information may profit particularly much from a chance to discover their erroneous decision mechanism through increased feedback. On the other hand, such individuals may be particularly overoptimistic regarding their ability to improve their trading outcomes via smart strategy adjustments and thus be particularly resistant to feedback. The experiment's overlapping treatment design is capable of addressing these concerns. Apart from experimental groups receiving the two respective treatments individually, a further experimental group receives both treatments. This allows to detect potential overlaps of two individual effects and check for potential interactions between the two.

The considered interventions lead to large and significant changes in behaviour. The Detailed Feedback and Information Filter Treatment achieve reductions in trading activity of 13% and 17%, respectively. These results are consistent with the explanation of both inexperience as well as mistakes in information processing mistakes being drivers in excessive trading. Accordingly, these changes in behaviour are accompanied by the expected changes in net returns, with the Detailed Feedback and Information Filter Treatments leading to sizeable increases in net returns of 0.42 and 0.33 percentage points. Moreover, the combined treatment shows virtually no interaction effects between the two individual treatments in the behaviour of the participants. Thus, these findings point towards both channels of excessive trading being independent drivers in creating the phenomenon, acting alongside each other.

These results provide new insights regardig the value of information in individual stock-trading. In an interactive market setup, Huber (2007) finds that average-informed traders perform the most poorly, while trading profits significantly increase with decreasing as well as increasing levels of information. The returns on information in such a setting may thus be non-monotonic. My results corroborate the view that traders may do better when receiving richer as well as poorer information. Crucially, however, the two interventions concern different *types* of information. Considering that their effects stack, my results suggest that distinct informational stimuli affect decision-making via independent channels rather than people acting on their overall informedness.

Moreover, I use personal information from the experiment's follow-up questionnaire to investigate the channels of overtrading. My analysis confirms previous studies (Biais et al., 2005; Grinblatt and Keloharju, 2009) that identify overconfidence and gambling activity as major drivers. On the other hand, contrary to Barber and Odean (2001) I do not find any link between a participant's gender and their level of overconfidence. In fact, male participants tend to trade less than women. Also, the treatment effects on participants' trading profits do not appear to arise from encouraging more risky portfolios as in Gneezy and Potters (1997). Instead, my results seem to stem from a novel channel of information management on trading outcomes.

At a practical level, my findings may be of interest to private individuals and professionals alike. Since the discussed experimental interventions solely concern which information should be presented when, they can easily be imitated on home computers with limited technical proficiency. These interventions could be implemented decentrally by households themselves or even offered as tools and applications by trading platforms or financial regulatory bodies acting in the public interest. Linnainmaa et al. (2018) further show that professional financial advisors may be subject to the same biases and mistakes as laymen, leading to high-cost overtrading both on behalf of their clients as well as within their personal portfolios, leaving room for my proposed interventions to exert benefits.

This article is organised as follows. The remainder of this section reviews related literature. Section 2.2 discusses the experimental design, provides a descriptive analysis of the data and verifies the success of the experiment's randomisation. Section 2.3 presents the experiment's main results, while Section 2.4 discusses possible channels. Section 2.5 concludes.

2.1.1 Related Literature

Tobin (1978) famously proposed to counter excessive trading volumes by introducing a financial transaction tax. Although his recommendation was targeted towards international money markets, Keynes (1936) argued for the same in the domain of stock trading, receiving further theoretical foundation from Palley (1999).

This paper, on the other hand, attempts to achieve the same effect solely by deliberately choosing which information to supply to and which information to withhold from the trader. In doing so, it aims at averting losses incurred from overtrading and thus increase profits. Fama (1970)'s hypotheses on information efficiency provide a benchmark for the impact of information on stock trading outcomes. He states that prices 'fully reflect all available information at all times' (p. 385). Therefore, all variations of additional information should not have an effect on trading behaviour.

Findings from numerous studies dispute this hypothesis. In zero-sum market games, better-informed individuals can profit by exploiting worse-informed individuals, showing that gathering information can, indeed, incur benefits (Copeland and Friedman, 1992). Importantly, such gains need not be monotonic in the level of information. In an experimental setup with multiple levels of informedness, Huber et al. (2008) show that only the best-informed traders can profit from their informational advantage, while in Huber (2007) average-informed traders obtain significantly lower returns than the worst-informed. On financial markets, less information can be better.

However, as these studies were carried out in zero-sum settings where trading gains are shifted across traders, their predictions have no bite in contexts where individuals trade in a non-interactive manner as price-takers on financial markets. Blackwell (1951) provides a well-known theoretical argument for individual decision-making, claiming generally that a rational prediction procedure can never perform better with less information. But these results need not hold true if individuals incorporate information incorrectly into their decision-making. Recent studies have accumulated rich evidence suggesting that this may indeed be true for individuals on the stock market. Frydman and Wang (2019), making use of a natural experiment, show that changes in stock price salience, while leaving available information unchanged, increases investors' tendencies to sell winning stocks and hold on to losers. Biais et al. (2005) provide evidence from a computer laboratory experiment, linking participants' overconfidence with their trading behaviour.

Given these insights, surprisingly little research has been done on how systematic changes in information management may affect trading behaviour in non-interactive contexts. As an early example one may consider Benartzi and Thaler (1995), showing theoretically that if stock traders are both myopic and loss-averse, they tend to select excessively non-risky portfolios with detrimental effects to their trading profits. Conducting a direct experimental test of the theory, Gneezy and Potters (1997) show that this tendency can be reversed by decreasing the frequency in which subjects receive performance updates on their allotted portfolios. Through the channel of more risky portfolios, restricted access to information leads to significant increases in returns. Andries and Haddad (2019) further argue theoretically that individuals may benefit from deliberately restricting their access to information for purely affective reasons. Loss averse traders, despite losing out on returns, may enjoy net-utility benefits from less frequent updates on their portfolio performances. With this article, I hope to contribute to this discussion.

2.2 Experimental Design

In order to study the impact of varying information sets on trading outcomes in stock markets, I conduct a computer laboratory experiment. All experimental treatments will merely vary the information which participants receive while leaving everything else unchanged. I will therefore begin by explaining the commonalities throughout all of the experiment's iterations and then describe how the individual treatments differ in design.

Participants trade on the basis of real-life stock data, in particular monthly returns from stocks listed on Standard and Poor's 500 Index from January 1960 to July 2017 (all data taken from the Bloomberg Terminal). At the beginning of each round, participants are randomly allotted a random stock in a random month from the basket. For each of these rounds, they must then make a binary stylised speculative trading decision: they either hold their stock or trade it. I choose data from real-life stocks as the basis for trading in the experiment in an attempt to ensure that insights from the computer laboratory remain valid outside of it. In particular, there exists ample evidence that excessive trading is linked with mistakes in expectation formation (i.e. overconfidence). It is unclear to what degree such distortions would carry over to different data generating processes, such as synthetic stock valuations drawn from pre-set distributions.[2]

If a participant decides to hold the stock, she will simply receive the returns of her allotted stock for the allotted month. If, on the other hand, she decides to trade, she will receive the returns for another random stock at a random month, while a fixed commis-

[2]For practical considerations, I further randomise the allocation of stocks in every period on the individual level as opposed to, say, randomly determining a stock which all participants hold in a given period. Monthly stock returns are highly volatile and a single outlier might otherwise have a large impact on both the participants' choice behaviour as well as their profits from trading. Randomising provides some robustness towards such concerns, while not introducing a systematic bias towards any of the experimental groups.

sion fee of five percentage points will be deducted from her earnings for the round. The net profits, denominated in percent of return, are then added to (or subtracted from) the participant's total earnings at the end of each round. At the end of the experiment, these earnings are converted and paid out at a rate of $1\% = \pounds 0.08$. Participants receive an additional fix reward of £6 for completing an experiment. Overall earnings are capped from below at £6 and from above at £20. The average final payment for the experiment was £6.90, which is roughly in line with the UK minimum wage for young adults at the time the experiment was conducted. Participants who did well in the experiment had the ability to substantially increase their payoffs. The standard deviation for final payoffs is £2.18 and sereval participants earned more than £10 from the experiment, providing ample incentives to engage with the experiment in trying to maximise monetary payoffs. Moreover, as seen in Table 2.7, the participants' choice behaviour is not significantly driven by their initial level of wealth for the most favoured model specification (significance further drops strongly when outliers are dropped). The experiment's monetary incentive effects therefore do not appear to abate with increasing wealth of the participants.

In order to make this decision, participants in the control group receive information on the past performance of their allocated stock. Specifically, they are shown a graph of the returns the stock has yielded in each of the past 10 months (see Figure C.7 in the appendix for an example). There are 30 rounds in the experiment and, up to the accumulation of earnings across periods, all rounds are technically independent from each other so that nothing the participant does in any given round may affect what will happen in another. Also, all participants trade independently of another, so that they may not affect each other.

As illustrated in Figure 2.2, there are two treatments in the experiment implemented in a 2×2 design. All experimental groups are played in each session, where participants are randomly allocated according to the four quadrants in Figure 2.2 and remain in their given treatment for the entirety of the session. Following the main experiment, all participants are asked to answer a short, non-incentivised follow-up questionnaire. Personal characteristics elicited from the survey will be used to supplement the empirical analysis (see Section 2.4).

In total, I conducted 18 sessions with a total of 352 participants. All sessions were held in the University of Glasgow, with 15 sessions from October 2017 to February 2018 and 3 further sessions in February 2019. On average, sessions lasted about 60 minutes. The experiment was fully computerised using z-Tree of Fischbacher (2007).

Figure 2.2: 2×2 experimental design.

Control	Detailed Feedback (Treatment I)
Information-Filter (Treatment II)	Treatment I & II

2.2.1 Experimental Treatments

Treatment I: Detailed Feedback

In this treatment, participants receive more information than the control group. In particular, participants will receive detailed feedback after each round regarding their performance (see Figure C.8 for an exemplary screen). Investors receive information on the gross performance of their chosen option and the resulting net profits. Also, they are told how much they would have earned if they would have chosen differently, including a split of net profit posts. Finally, they are prompted with an explicit evaluation whether, in the given instance, they have chosen the most profitable option or not.

Treatment II: Information Filter

In this treatment, participants receive less information than in the control group. In particular, participants are subject to an algorithmic information filter which determines whether they will receive information on their allotted stock's performance history.

The discrimination on whether the participants receive this information is implemented as follows: Each round, a random natural number z from 0 to 9 is drawn with uniform probabilities. The algorithm then counts the successive run of either positive or negative returns of the participant's allotted stock for the given round, starting with the most recent return. Let θ denote this number. For example, if the stock yielded positive returns for the last three months preceded by a negative return, $\theta = 3$. If conversely the stock yielded negative returns successively for the last three months preceded by a positive return, then also $\theta = 3$. If the stock yielded positive returns for the last month preceded by nine negative returns, then $\theta = 1$.

The algorithm then compares z and θ. If $\theta > z$, then the participant will indeed receive information on the stock's past performance and the participant will follow perfectly the protocol of the control group for the round. If the converse holds true, with $\theta \leq z$, then the participant will not receive any information on the allotted stock's performance, but will instead be told that information for the round has been restricted (see Figure C.8 for a screen display). The algorithm therefore tends to flag exceptions and suppress relatively moderate outcomes, subject to the random draw of z.

2.2.2 Data and Randomisation Check

Table 2.1 provides the means, standard deviation and minimum and maximum realisations of variables elicited in the experiment. Panel A lists variables on trading outcomes elicited during the main experiment. Net Returns denotes the participants' returns earned in a given period from stocks minus commission fees subject to trading activity. Gross Returns denotes the same outcome without deducting potential commission fees. Trading denotes the participant's choice in any given period to trade (coded with 1) or not (coded with 0). As can be seen, net returns are highly volatile, featuring a low absolute mean and a high standard deviation. This will make it difficult, although not impossible, to find statistically significant effects on this outcome due to the substantial noise.

Panel B in Table 2.1 lists select variables related to participants' personal characteristics, based on information in the experiment's follow-up questionnaire. I will now provide an account of how these measures are constructed, top to bottom. Screenshots of all the underlying questions are available in the appendix (see chapter C). The variable Male is coded with 1 if the participant reported to be male and 0 otherwise. Age simply repeats the stated age. I follow the methodology of Holt and Laury (2002) to elicit participants' risk sentiment. Participants have to choose ten times between two lotteries. For the Riskseeking Index, I count the number of risky choices and divide it by ten. For the Overconfidence Index, inspired by Biais et al. (2005), I use a version of the survey tool introduced in Alpert and Raiffa (1982). Participants are given ten questions such as: 'What is the height of the Mt. Everest (in meters)'. They are then instructed to give both a low and a high estimate, as close as possible to what they believe the true answer is but such that they are 90% certain that the true answer lies in the stated interval. To arrive at the index, I then count the number of times that the true answer did, in fact, lie outside the stated interval and divide this number by the total amount of questions. For the Gambling Index, participants can state four levels of intensity of gambling activity in their weekly routines. I code these four levels in ascending intensity with 0, 0.33, 0.66 and 1. Both Stock Experience Basic and Stock Experience Advanced are dummy variables aimed at eliciting the level of participants' experience in trading stocks, based on simple 'yes' or 'no' questions. For the Irregular BMI variable, I first use stated height and weight in order to calculate the Body Mass Index according to the standard formula $\frac{\text{Weight(kg)}}{\text{Height(m)}^2}$. I then compare this number to the standardised benchmarks 18.5 and 25 to check whether a person classifies as underweight or overweight, respectively. Wealth states the self-reported sum of the value of all assets owned by the participant.

Table 2.1: Descriptive Statistics

	Mean	Std. Dev.	Min.	Max.
	(1)	(2)	(3)	(4)
Panel A: Trading Outcomes				
Net Returns	-0.24	10.32	-79.97	171.78
Gross Returns	1.42	10.05	-79.97	176.78
Trading	0.33	0.47	0	1
Exp. Payment	6.9	2.18	6	20
Panel B: Individual Characteristics				
Male	0.40	0.49	0	1
Age	24.77	7.54	17	69
Riskseek. Index	0.50	0.17	0	1
Overconf. Index	0.72	0.16	0.1	1
Gambling Index	0.17	0.26	0	1
Stock Exper. Basic	0.20	0.40	0	1
Stock Exper. Adv.	0.13	0.34	0	1
Low Self Est. Index	0.63	0.17	0.06	1
Irregular BMI	0.26	0.44	0	1

Notes: This table provides descriptive statistics within the sample for some key variables. 'Trading Outcomes' are noted at the decision-level (30 per participant), while 'Individual Characteristics' are stated at the participant-level.

Table 2.2: Randomisation Check, Part 1: Means by Randomisation Groups

	Top Left (1)	Top Right (2)	Bottom Left (3)	Bottom Right (4)
Male	0.43	0.39	0.37	0.44
	(0.50)	(0.49)	(0.49)	(0.50)
Age	24.25	25.44	24.84	24.55
	(5.09)	(9.86)	(7.07)	(7.44)
RiskSeek. Index	0.47	0.50	0.52	0.50
	(0.17)	(0.18)	(0.17)	(0.16)
Gambling Index	0.19	0.11	0.17	0.18
	(0.28)	(0.22)	(0.28)	(0.26)
Overconf. Index	0.70	0.73	0.73	0.73
	(0.17)	(0.16)	(0.14)	(0.17)
Stock Exper. Basic	0.21	0.22	0.22	0.16
	(0.41)	(0.41)	(0.41)	(0.37)
Stock Exper. Adv.	0.16	0.12	0.14	0.09
	(0.37)	(0.33)	(0.35)	(0.29)
Low Self Est. Index	0.65	0.63	0.62	0.62
	(0.15)	(0.18)	(0.17)	(0.17)
Irregular BMI	0.27	0.25	0.22	0.30
	(0.44)	(0.43)	(0.41)	(0.46)

Notes: This table provides the means (with standard deviations in parentheses) of some key personal characteristics for each of the experiment's four randomisation groups (see Figure 2.2).

I use information regarding the participants' personal characteristics to check whether there are significant differences between the samples of the randomisation groups (the four quadrants in Figure 2.2). This verifies the succes of the experiment's randomisation. Table 2.2 displays the means and standard deviations of some personal characteristics for the four subsamples. Based on these, Table 2.3 checks for significant differences across all combinations of the subgroups. I find that out of 60 comparisons, only two are significant at the 10% level, which is consistent with chance.

Table 2.3: Randomisation Check, Part 2: Differences across Randomisation Groups

	Col. 1-2 (1)	Col. 1-3 (2)	Col. 1-4 (3)	Col. 2-3 (4)	Col. 2-4 (5)	Col. 3-4 (6)
Male	0.033	0.056	-0.008	0.022	-0.042	-0.064
	(0.074)	(0.073)	(0.075)	(0.073)	(0.075)	(0.074)
Age	-1.191	-0.584	-0.294	0.606	0.896	0.289
	(1.176)	(0.924)	(0.963)	(1.286)	(1.329)	(1.100)
Riskseek. Index	-0.029	-0.041	-0.023	-0.012	0.005	0.017
	(0.027)	(0.026)	(0.025)	(0.026)	(0.026)	(0.025)
Gambling Index	0.074*	0.014	0.006	-0.059	-0.068*	-0.008
	(0.038)	(0.042)	(0.042)	(0.038)	(0.037)	(0.042)
Overconf. Index	-0.034	-0.038	-0.035	-0.003	-0.000	0.002
	(0.025)	(0.024)	(0.026)	(0.023)	(0.026)	(0.024)
Stock Exper. Basic	-0.011	-0.011	0.048	0	0.060	0.060
	(0.062)	(0.062)	(0.059)	(0.062)	(0.060)	(0.060)
Stock Exper. Adv.	0.044	0.022	0.074	-0.022	0.029	0.051
	(0.053)	(0.054)	(0.051)	(0.051)	(0.047)	(0.049)
Low Self Est. Index	0.020	0.033	0.031	0.012	0.011	-0.001
	(0.025)	(0.025)	(0.025)	(0.027)	(0.026)	(0.026)
Irregular BMI	0.016	0.051	-0.027	0.034	-0.044	-0.079
	(0.067)	(0.065)	(0.070)	(0.034)	(0.069)	(0.068)

Notes: This table provides a randomisation check for the experiment. It shows the differences between the means of participant characteristics denoted in Table 2.2 (standard errors in parantheses). Alternative hypothesis is the two-sided inequality. Statistical significance denoted by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2.3 Results

In this section I will examine the effectiveness of the experimental treatments in curbing excessive trading. I begin by introducing the analysis model and and overall results and then discuss them with respect to the experiment's treatments. The most granular data, meaning trading choice per participant per period, will represent a unit of observation. Letting i index participants, j sessions and t experimental periods, the regression equation can thus be written as follows:

$$\begin{aligned} TRADING_{ijt} = & \beta_0 + \beta_1 \cdot FEEDBACK_{ij} + \beta_2 \cdot FILTER_{ij} \\ & + \beta_3 \cdot INTERACT_{ij} + \phi_j + \tau_t + \varepsilon_{ijt} \end{aligned} \tag{2.1}$$

where $TRADING_{ijt}$ denotes a participant's choice for a given period, coded 0 if the participant decided to hold and 1 if the participant decided to trade. $FEEDBACK_{ij}$ and $FILTER_{ij}$ are dummy variables for the respective experimental treatments and $INTERACT_{ij}$ is a dummy variable capturing potential interaction effects for participants who receive both experimental treatments. ϕ_j and τ_t are dummy vectors controlling for potential session- and round-fixed effects which will be added and omitted from the analysis for robustness. The error term is ε_{ijt}. Table 2.4 displays the respective regression coefficients from a pooled OLS-regression with robust standard errors. I run an identical regression with a net returns, measuring a participant's gross returns in a given from holding stocks minus any potential commission fees, as the dependant variable. Table 2.5 displays the results.

Table 2.4: Effect of Experimental Treatments on Trading

	(1) Trading	(2) Trading	(3) Trading	(4) Trading
$FEEDBACK$	-0.0554***	-0.0550***	-0.0554***	-0.0550***
	(-4.20)	(-4.16)	(-4.21)	(-4.16)
$FILTER$	-0.0734***	-0.0717***	-0.0734***	-0.0717***
	(-5.60)	(-5.47)	(-5.60)	(-5.47)
$INTERACT$	0.00438	0.00470	0.00438	0.00470
	(0.24)	(0.26)	(0.24)	(0.26)
Fixed Effects				
Session		✓		✓
Period			✓	✓
N	10560	10560	10560	10560

Notes: This table provides the coefficient estimates of the experimental treatments on the participant's trading choices. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.5: Treatment Effects on Net Returns

	(1) Net Returns	(2) Net Returns	(3) Net Returns	(4) Net Returns
$FEEDBACK$	0.448*	0.432*	0.447*	0.432*
	(1.92)	(1.84)	(1.92)	(1.84)
$FILTER$	0.357	0.338	0.357	0.337
	(1.55)	(1.34)	(1.55)	(1.46)
$INTERACT$	-0.232	-0.205	-0.230	-0.204
	(-0.71)	(-0.63)	(-0.70)	(-0.62)
Fixed Effects				
Session		✓		✓
Period			✓	✓
N	10350	10350	10350	10350

Notes: This table provides the coefficient estimates of the experimental treatments on net returns. Outliers beyond the 1st- and 99th-percentile in the net return-distribution have been dropped. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.3.1 Detailed Feedback Treatment

As can be seen in Table 2.4, the Feedback Treatment leads to a highly significant reduction in the trading behaviour of the participants. Compared to the Control Group, participants whose propensity to generate experience from their past actions are bolstered by the treatment are about 13% less likely to trade, but instead adopt a more passive strategy of holding their allocated stocks. Moreover, Table 2.5 shows that these changes in behaviour lead to the expected changes in the participants' payoff-relevant metric. The treatment increases monthly net returns by a sizeable 0.45 percentage points, significant at the 10%-level.

This set of findings corroborates the explanation proposed by Seru et al. (2009) for the excessive trading puzzle in two accounts. First, my results confirm that a higher level of experience leads to increases in the trading performance of individual investors. Second, my results confirm the type of learning responsible for these gains. Rather than increasing the participants' capabilities of spotting advantageous opportunities, the richer availability of performance feedback allows traders to discover their own ineptitude of doing so. Inflated misconceptions on the own ability of generating profits from active trading may thus be discovered and addressed by an increased access to information.

These results overall may thus reflect a rational learning process, rendering excessive trading as the outgrowth of uncertainty of the traders regarding themselves.

2.3.2 Information Filter Treatment

Table 2.4 shows that reducing access to information regarding the allocated stock's performance history in the Information Filter Treatment leads to significant reductions of excessive trading. Participants who do not possess this information to estimate the likelihood of the stock generating bad performances in the upcoming period predominantly hold the stock instead of trying to trade out of their losses. Compared to the Control Group this accounts for a decrease of 17% of trading activity.

If participants do not have any information to project the performance of their allocated stock, refraining from trading and thus avoiding the commission fee is, in fact, unconditionally the correct choice. This behaviour is therefore consistent with a rational trader attempting to maximise monetary payoffs. If trading activity in the presence of information was, however, motivated by calculated speculative trading, reducing access to information should not lead to an average increase in payoffs. Table 2.5 shows that, in direction, the Information Filter Treatment indeed does lead to a large increase in payoffs, although this effect misses out on significance at the 10%-level due to high volatility in the dependant variable. A stronger picture crystallises, if focus is restricted on rounds where the historic price information has been indeed suppressed. In Table 2.6, INFO RESTRICTION is a dummy variable switching on for any period in which participants did not receive such information. As can be seen, Net Returns in such periods is then about 0.30 percentage points higher compared to all other rounds, marking a significant increase at the 10%-level.

These findings are thus consistent with excessive trading being the result of mistakes in decision-making. Distorted perceptions of perceived information, caused for example by overconfidence in the precision of one's own ability to infer valuable signals from acquired information, can lead to overreactions to such information and, thus, overtrading. If such overreactions are sufficiently prevalent, reducing access to information may improve the average quality of trading decisions rather then decreasing it.

Table 2.6: Effect of Information Restriction on Net Returns

	(1) Net Returns	(2) Net Returns	(3) Net Returns	(4) Net Returns
INFO RESTRICTION	0.302*	0.291*	0.304*	0.293*
	(1.81)	(1.74)	(1.82)	(1.75)
Fixed Effects				
Session		✓		✓
Period			✓	✓
N	10320	10320	10320	10320

Notes: This table provides the coefficient estimates of information restriction on net returns. The independent variable is a dummy, switching on if information on the historic returns (subject to the algorithm in the information filter treatment) has indeed been restricted. Unit of observation is the decision. Outliers beyond the 1st- and 99th-percentile in the net return-distribution have been dropped. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.3.3 Interaction Treatment

Due to the 2×2 experimental design (see Figure 2.2), a segment of the participant pool receives both the Detailed Feedback and the Information Filter Treatment, named the Interaction Treatment. For this group, analytically both the effect of the first treatment, the second treatment as well as potential interaction effects between the two treatments are considered in determining behaviour and outcomes. Using the terminology of Equation (2.1), the total effect is represented by $\beta_1 + \beta_2 + \beta_3$.

Table 2.4 reveals that the effects of the individual treatments stack nearly seamlessly in effectiveness for curbing excessive trading in the group. Both the regression coefficient capturing possible interaction effects as well as the respective t-statistic are very close to 0. The evidence therefore neither suggests that the effects of the individual treatments reinforce each other nor that they oppose each other. Instead, both lack of knowledge regarding one's own capability as well as mistakes in incorporating new information into decision-making appear to be independent channels in driving excessive trading. Consequently, participants receive the full benefits of both exogenously increased experience as well as concealment of potentially harmful information, leading to dramatic changes in behaviour.

These changes are generally reflected in the group's net returns, too. Table 2.5 shows that there are some negative interaction effects between the experiment's two individual treatments, but these are statistically highly insignificant. In any case, this interaction effect is lower than any of the two individual treatment effects, so that the group receiving both treatment earns the highest net returns on average (difference of the combined effect to the control group significant at the 5%-level with p-value 0.0174). It should be stressed that, apart from their significance, these effects are large. For comparison, the mean net returns for participants receiving none of the treatments, excluding outliers beyond the 1st- and 99th-percentile, lay at -0.69%, while the group receiving both treatments enjoys a 0.56 percentage point increase in net returns. Both interventions together are thus able to reverse more than 80% of the losses incurred from overtrading.

2.4 Discussion

2.4.1 Overconfidence and Gambling

Although the phenomenon of overtrading amongst private individuals, *i.e.* high trading volumes leading to subpar net returns, is empirically well-established, its causes are subject to debate. A strand of literature contends that the observed behaviour is the result of overstatement of one's own ability to infer information from any given data, thus leading to overprecise beliefs. This hypothesis dates back at least to Bondt and Thaler (1985), who show that actors in the stock market overreact to price movements. Biais et al. (2005) support this view with experimental evidence, showing that overtrading is linked to a particular of a psychological disposition towards overreacting dubbed 'overconfidence': when predicting future events, people do not sufficiently account for randomness and overweigh their private information.

While Grinblatt and Keloharju (2009) verify this result with real-life data, they propose a second, independent driver for overtrading. The authors find that trading activity correlates at the individual level with the amount of speeding tickets received. As an explanation, they argue that both overtrading and speeding may be the result of *sensation seeking*. Stock trading, in this case, assmumes the role of a thrilling gambling activity. Gao and Lin (2014) further corroborate this hypothesis with findings from natural experiments, showing that individuals partially substitute stock trading with playing the lottery. In this case, people may 'consume' overtrading, trading off its leisure value with pecuniary drawbacks.

Table 2.7: Effects of Personal Characteristics on Trading Choices.

	(1) Trading	(2) Trading	(3) Trading	(4) Trading
Male	-0.0201**	-0.0245**	-0.0201**	-0.0245**
	(-2.01)	(-2.40)	(-2.01)	(-2.40)
Age	0.00163**	0.00197***	0.00163**	0.00197***
	(2.28)	(2.62)	(2.28)	(2.62)
Riskseek. Index	-0.0268	-0.0285	-0.0268	-0.0285
	(-0.99)	(-1.02)	(-0.99)	(-1.02)
Gambling Index	0.0582***	0.0625***	0.0582***	0.0625***
	(3.21)	(3.32)	(3.20)	(3.32)
Stock Exper. Basic	-0.0275	-0.0255	-0.0275	-0.0255
	(-1.63)	(-1.48)	(-1.63)	(-1.48)
Stock Exper. Adv.	-0.00830	-0.00549	-0.00830	-0.00549
	(-0.43)	(-0.28)	(-0.43)	(-0.28)
Low Self Est. Index	0.00693	0.0230	0.00693	0.0230
	(0.25)	(0.80)	(0.25)	(0.81)
Overconf. Index	0.102***	0.101***	0.102***	0.101***
	(3.61)	(3.41)	(3.62)	(3.42)
Irregular BMI	0.0142	0.0120	0.0142	0.0120
	(1.30)	(1.08)	(1.30)	(1.08)
Wealth	-5.43e-08*	-4.53e-08	-5.43e-08*	-4.53e-08
	(-1.96)	(-1.61)	(-1.96)	(-1.61)
Fixed Effects				
Session		✓		✓
Period			✓	✓
Observations	10200	10200	10200	10200

Notes: This table provides the coefficient estimates of participants' personal characteristics on their trading choices. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In this section, I use information from the experiment's follow-up questionnaire to contribute to this discussion. Table 2.7 depicts the extent to which personal characteristics determines the participants' trading behaviour (see Section 2.2.2 for a detailed description of the construction of all variables). The results support both strands of literature. In all specifications, an individual's gambling frequency and overconfidence are both the most consistent (significant at the 1%-level) as well as the strongest (evaluated at a change of 1 standard deviation in the all characteristics) predictors of overtrading.

The results, however, contrast with findings from Barber and Odean (2001). Several studies show gender-differences in different domains of economic decision-making (see Croson and Gneezy, 2009, for a review). Leaning on such evidence, Barber and Odean (2001) adopt a trader's gender as a proxy for the tendency to act overconfidently. They then show that men more than women tend to trade excessively. My results, on the other hand, do not show any systematic differences between men and women with respect to overconfidence (p-value for two-sided t-test is 0.42). Indeed, as seen in Table 2.7, male participants trade less than female ones.

A number of other studies have found divergent results regarding gender-differences in financial domains, suggesting a more complex underlying mechanism than a binary gender distinction. Although Deaves et al. (2008) find that in an experimental setting overconfidence increases trading activity, they neither find gender-differences in overconfidence and only limited support that males trade more often than females. Chen and Cheng (2018) find, using transactional data from the Taiwanese futures market, that men trade more often than women but also lose less in doing so, thus disputing an underlying difference in overconfidence as the driver. They further argue that gender differences in economic preferences may be driven by the contemporary level of gender-equality. In particular, a 'lack' of overconfidence may be driven by a lack of inequality. Evidence from Cho (2017) on math-tests across different countries support this hypothesis.

It should be mentioned, though, that of the discussed personal characteristics only gambling activity translates to a statistically significant effect on trading profits (see Table 2.8).

Table 2.8: Effects of Personal Characteristics on Net Returns

	(1) Net Returns	(2) Net Returns	(3) Net Returns	(4) Net Returns
Male	0.0852	0.0902	0.0878	0.0929
	(0.47)	(0.49)	(0.49)	(0.51)
Age	-0.00682	-0.00922	-0.00681	-0.00917
	(-0.57)	(-0.71)	(-0.57)	(-0.71)
Riskseek. Index	0.721	0.871*	0.723	0.873*
	(1.47)	(1.72)	(1.48)	(1.73)
Gambling Index	-0.561*	-0.588*	-0.562*	-0.589*
	(-1.77)	(-1.77)	(-1.77)	(-1.78)
Stock Exper. Basic	0.116	0.0852	0.119	0.0878
	(0.38)	(0.27)	(0.39)	(0.28)
Stock Exper. Adv.	0.367	0.308	0.361	0.301
	(1.07)	(0.88)	(1.05)	(0.86)
Low Self Est. Index	-0.258	-0.281	-0.254	-0.277
	(-0.51)	(-0.55)	(-0.51)	(-0.54)
Overconf. Index	0.123	0.120	0.115	0.112
	(0.24)	(0.22)	(0.22)	(0.21)
Irregular BMI	-0.296	-0.260	-0.293	-0.256
	(-1.52)	(-1.30)	(-1.50)	(-1.28)
Fixed Effects				
Session		✓		✓
Period			✓	✓
N	9992	9992	9992	9992

Notes: This table provides the coefficient estimates of participants' personal characteristics on their net returns. Outliers beyond the 1st- and 99th-percentile in the net return-distribution have been dropped. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.4.2 Robustness Towards Risk as a Channel

This article support the argument that information management can curb excessive trading in a non-interactive stock trading task as well as increasing the profits gained from it. These results could, however, simply be generated by a shift in the participants' risk adjustments. For example, traders could have a desire to hold low-risk portfolios and be willing to give up some expected returns in order to meet this goal. Therefore, if they observe that their allocated stock's returns have been highly volatile over the past months, they may rationally wish to swap this stock for an (expectedly) lower-risk, lower-payoff one. In turn, restricting access to such information would then implicitly foster risk-taking, decrease trading activity and increase expected returns.

Indeed, Gneezy and Potters (1997) find in a lab experimental setting that risk-taking and trading profits increase if feedback on performance information is supplied less frequently, while Larson et al. (2016) find in a natural experiment that informing traders less frequently on prices leads to the same qualitative behaviour. An unsuccesful replication by Beshears et al. (2016), on the other hand, calls this effect into question.

Risk considerations, however, do not appear to drive this paper's findings. In Table 2.9, Net Return Std. measures the volatility of participants' period-by-period net returns through their standard deviation (column (2) displays the same for gross returns). If the decreased trading volume and increased returns in the experimental treatments were a result of increased willingness to hold risky stocks, one would expect greater volatility in the associated payoff stream. However, for both specifications, allocation to any treatment group actually, if anything, lowers payoff stream volatility. Thus, in terms of a risk-reward trade-off, the trading performance assisted by the experiment's proposed information management dominates the alternative.

Table 2.9: Effects of Experimental Treatments on Portfolio Risk

	(1) Net Return Std.	(2) Gross Return Std.
FEEDBACK	-0.685	-0.600
	(-1.40)	(-1.20)
FILTER	-0.836*	-0.703
	(-1.67)	(-1.37)
INTERACTION	1.078*	0.926
	(1.70)	(1.43)
Fixed Effects		
Session	✓	✓
N	352	352

Notes: This table provides the coefficient estimates of the experimental treatments on portfolio risk. Unit of observation is the participant. Independent variables denote the standard deviation of the participants' returns earned throughout all periods; in col. (1) potential commission fees are deducted, in col. (2) not. Treatment variables are dummies, switching on if a participant belongs to the respective treatment. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Also, there is no evidence that participants attempt to trade out of risky stock allocations. In Table 2.10, the variable Hist Ret. Std. captures the volatility of the historical return of the participant's allocated stock. Naturally, the sample is restricted to periods where this information was not suppressed by the FILTER treatment and participants could indeed observe the stock's return history. Again, it can be seen that in both specifications higher observed volatility rather leads to less trading activity than more, although the effect size is very small. I therefore argue that the benefits from the experiment's information interventions arise from novel channels which are independent from those put forward by Gneezy and Potters (1997).

Table 2.10: Effect of Allocated Stock Risk on Trading Activity

	(1) Trading	(2) Trading	(3) Trading	(4) Trading
Hist. Ret. Std.	-0.000244*	-0.000234*	-0.000227	-0.000216
	(-1.73)	(-1.72)	(-1.50)	(-1.48)
Fixed Effects				
Session		✓		✓
Period			✓	✓
N	6346	6346	6346	6346

Notes: This table provides the coefficient estimates of historic risk on trading activity. Unit of observation is the decision. "Hist. Ret. Std." denotes the standard deviation of the historic returns of the allocated stock. The sample is restricted to decisions in which the participant could see this information. t statistics in parentheses. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.5 Conclusion

Excessive trading is a robust real-life phenomenon leading to economically large losses. In this paper, I present experimental evidence on functional counter-strategies. Both proposed measures lead to sizeable reductions in trading activity as well as increases in trading profits. Moreover, as the interventions are purely informational, these findings reveal new insights into the value of information in non-interactive stock trading by private individuals. Although I find evidence that providing performance feedback increases trading profits, perhaps even more surprising is that also restricting information on noisy predictors reduces trading activity.

The analysis further suggests that the results are not generated through increased risk-taking, a channel proposed in a similar setting by Gneezy and Potters (1997). Overall, the results may be seen as a synthesis of Seru et al. (2009)'s and Odean (1998b)'s hypotheses on the causes of excessive trading. The experiment's two individual treatments support both hypotheses. Combining both treatments reveals no interdependencies between them, but the effects are rather consistent with both hypotheses acting as independent channels contributing to the phenomenon.

Appendix C

Materials from the Experiment

Figure C.1: First instructions screen, common for all participants.

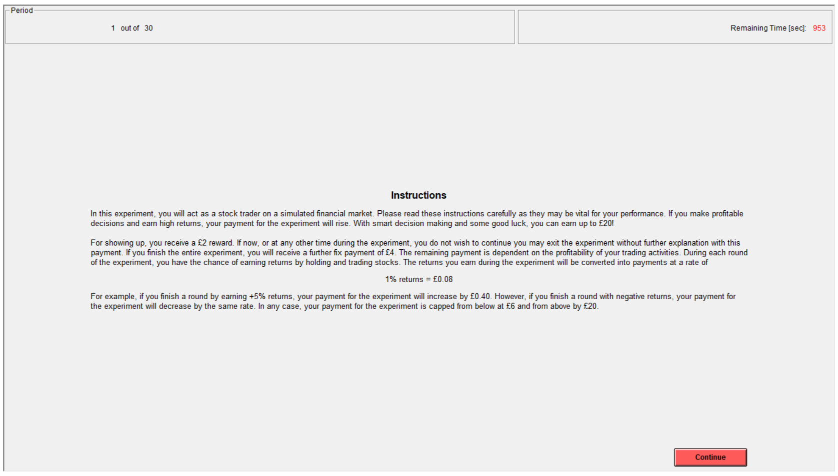

Figure C.2: Second instructions screen, common for all participants.

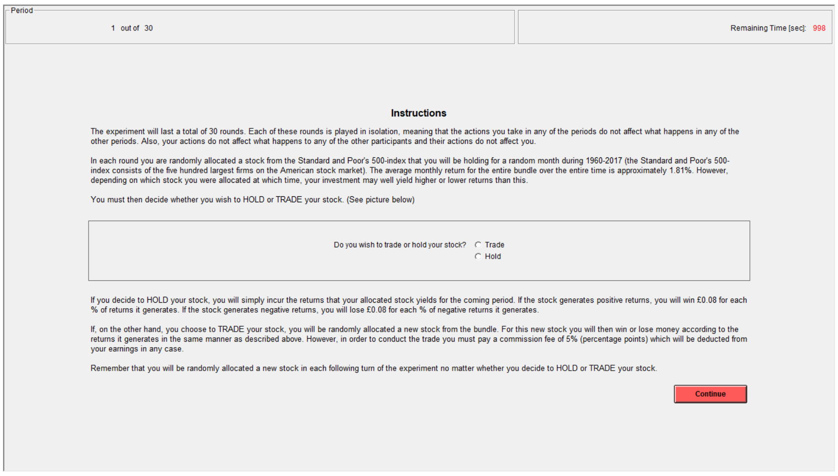

Figure C.3: Third instructions screen for Control Group.

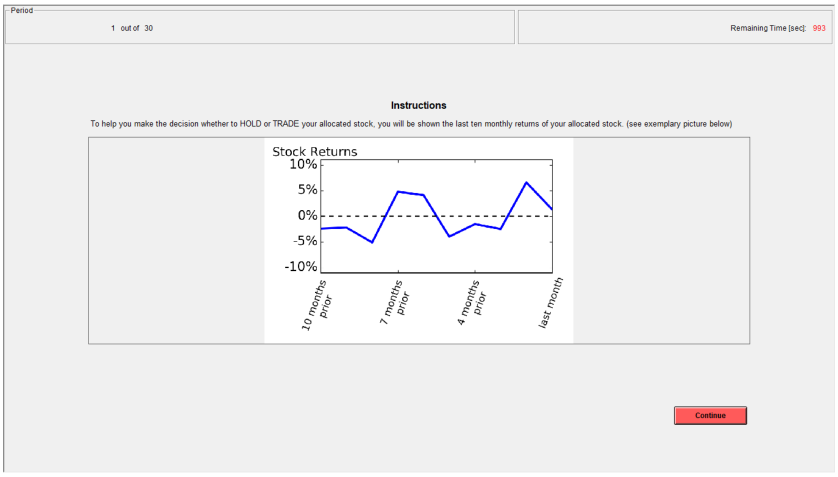

Figure C.4: Third instructions screen for Detailed Feedback Treatment.

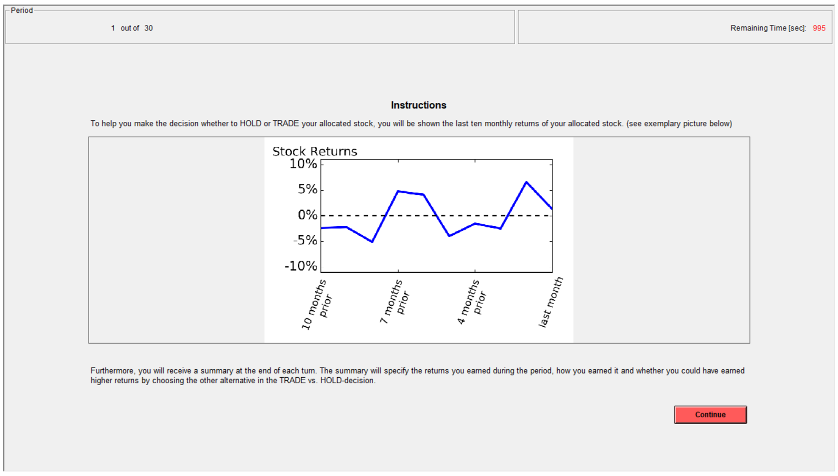

Figure C.5: Third instructions screen for Information Filter Treatment.

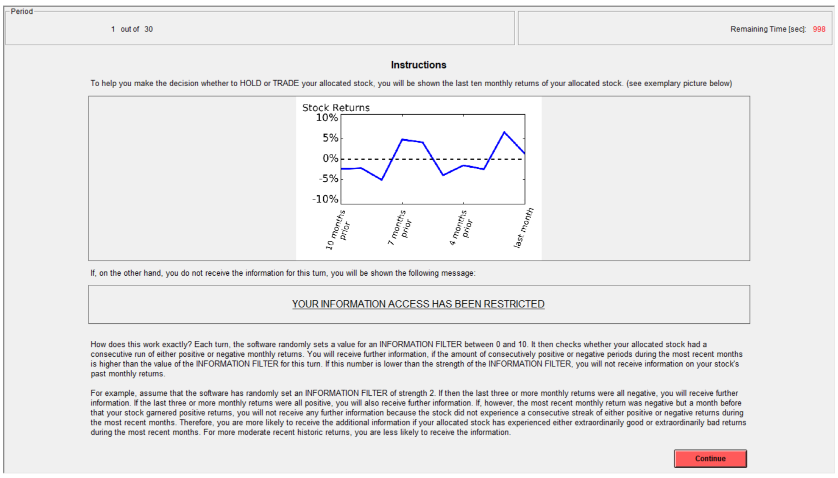

Figure C.6: Third instructions screen for Treatment Interaction Group.

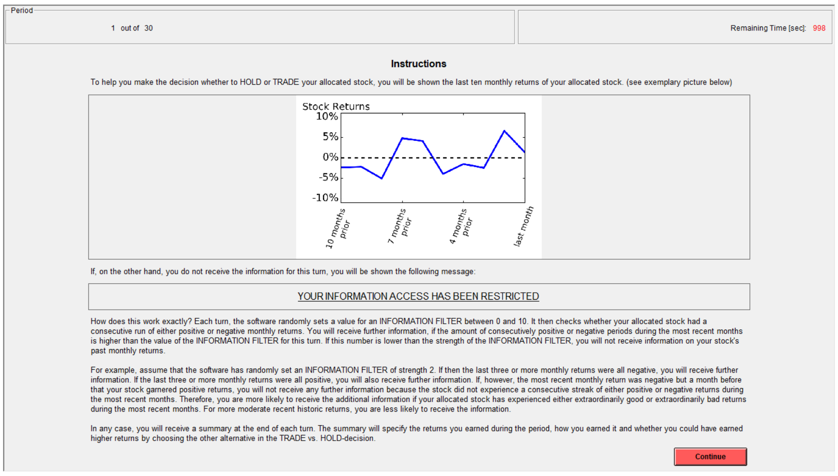

Figure C.7: Exemplary information of allotted stock's performance history in baseline group.

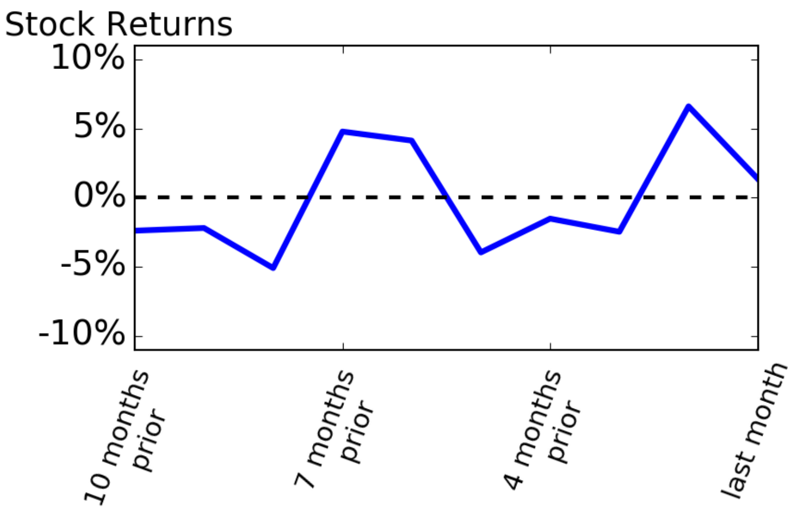

Figure C.8: Exemplary screen display for the Detailed Feedback Treatment.

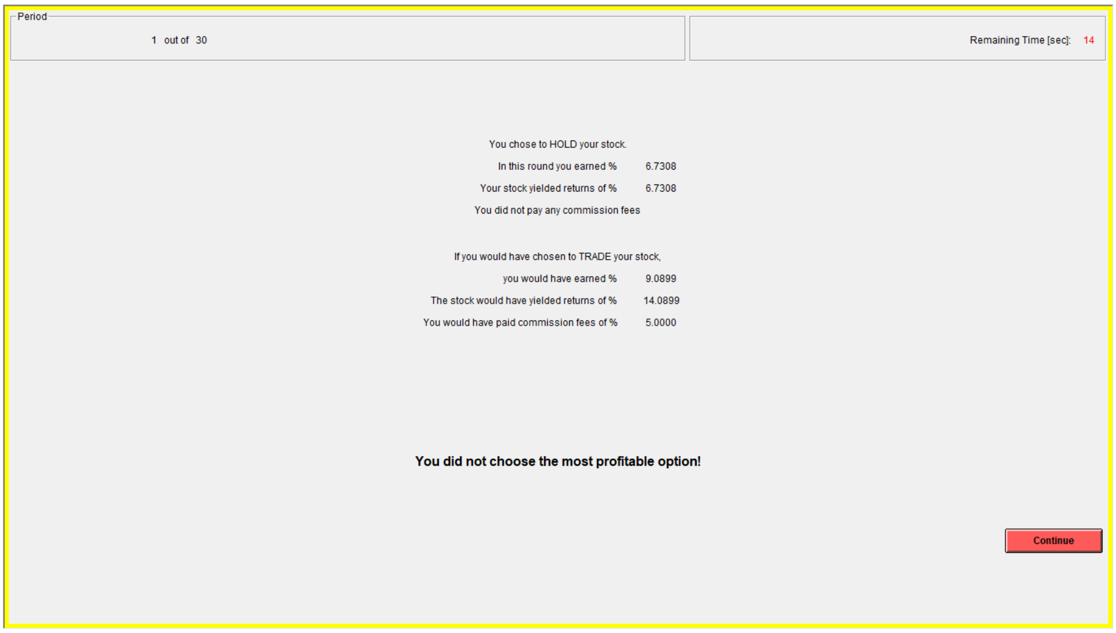

Figure C.9: Exemplary screen display for the Information Filter Treatment when information is restricted.

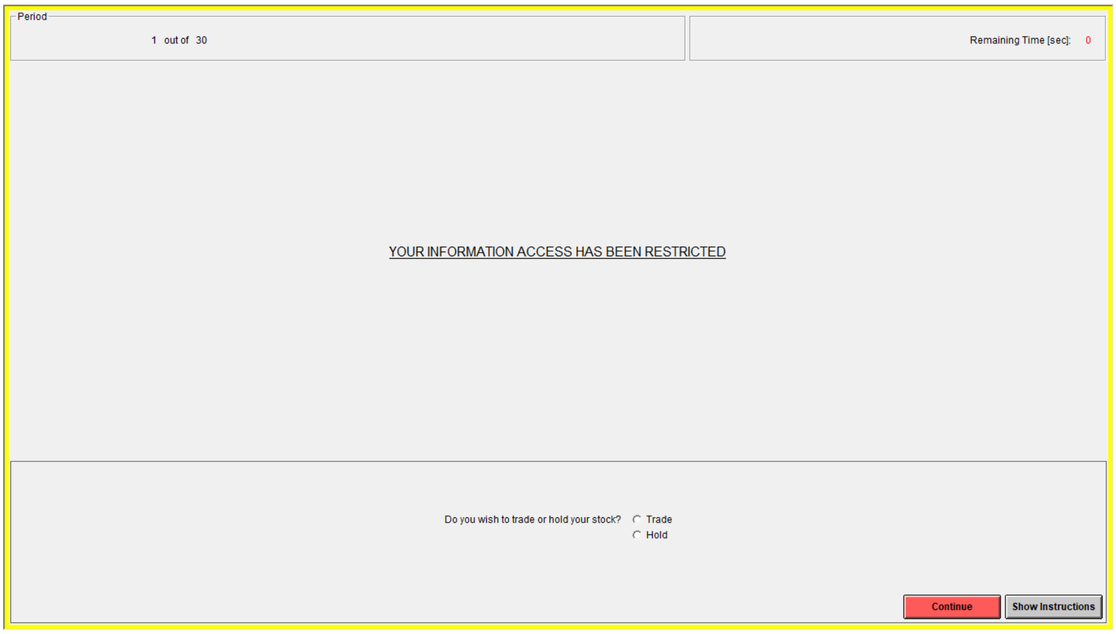

Figure C.10: Follow-up questionnaire, first screen display

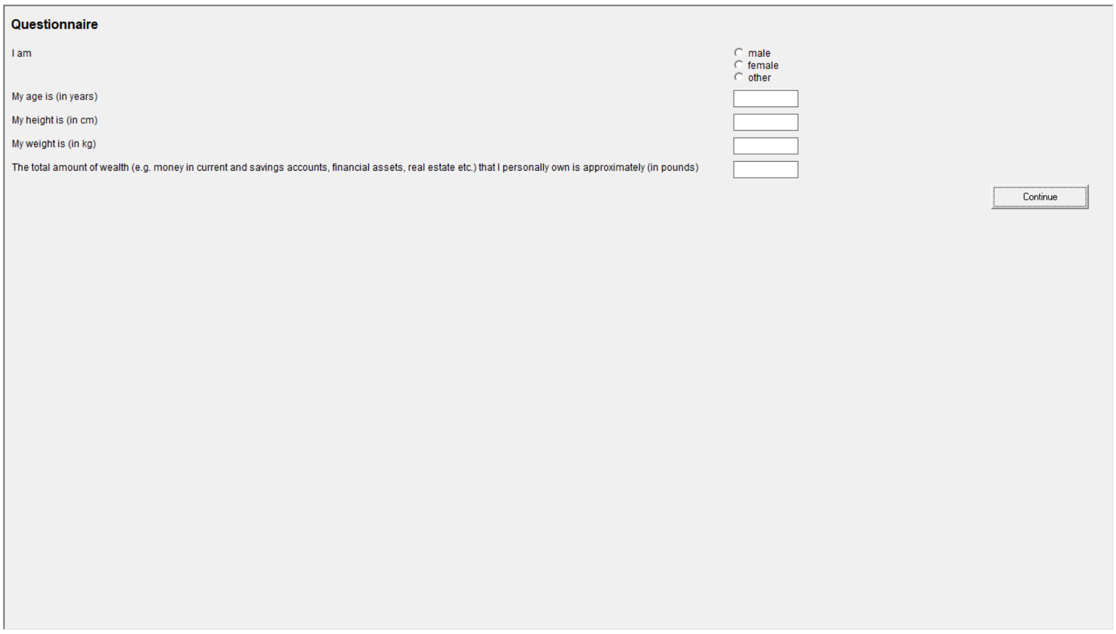

Figure C.11: Follow-up questionnaire, second screen display

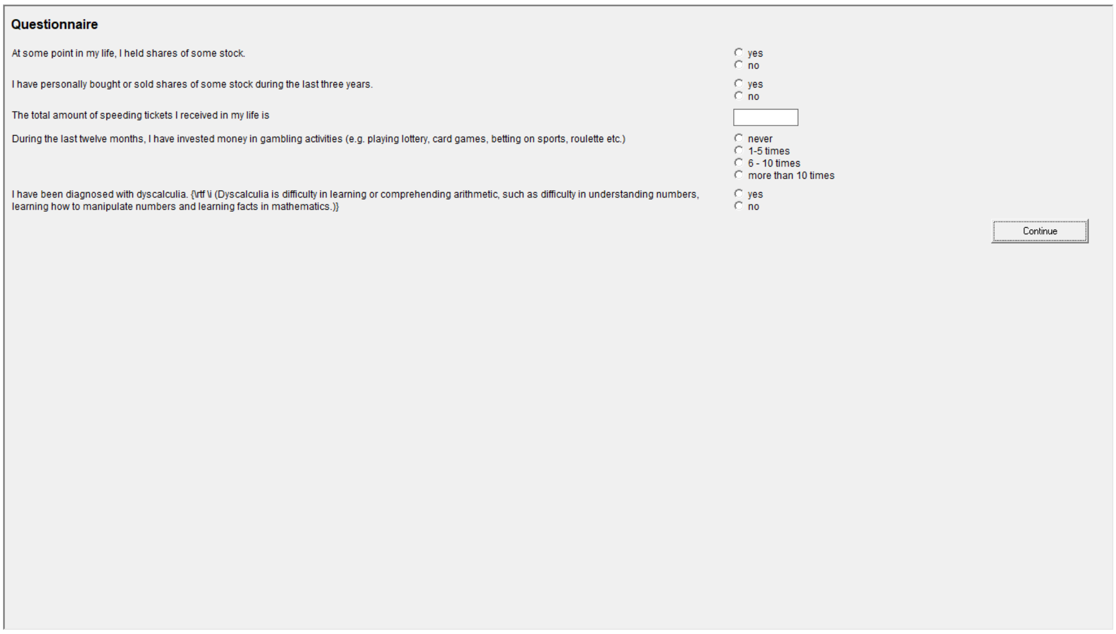

Figure C.12: Follow-up questionnaire, eliciting risk aversion

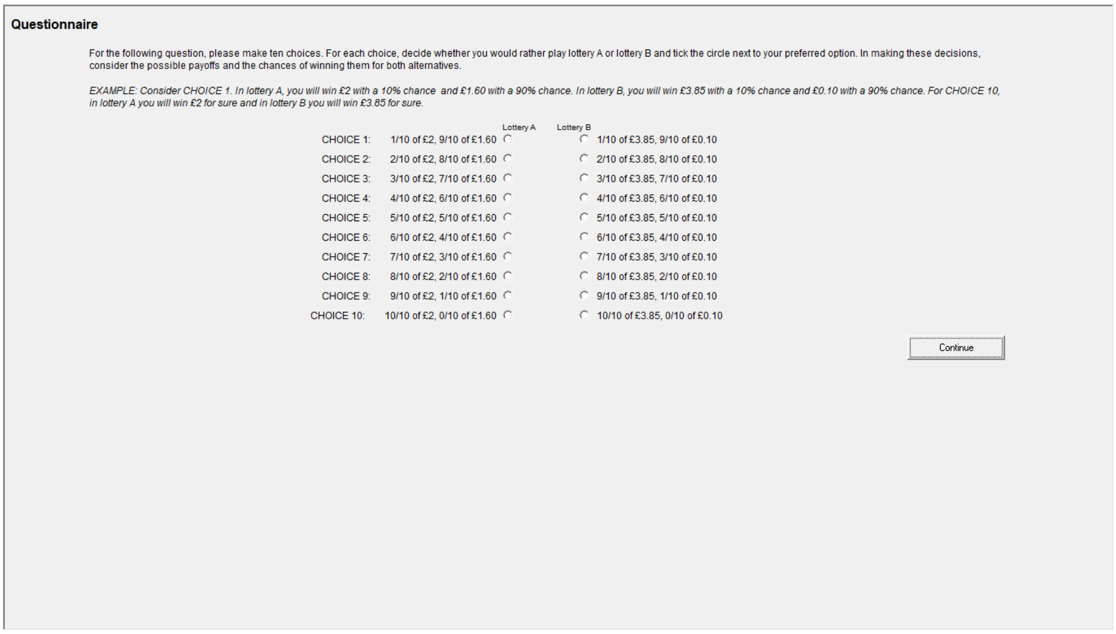

Figure C.13: Follow-up questionnaire, eliciting measure for overconfidence

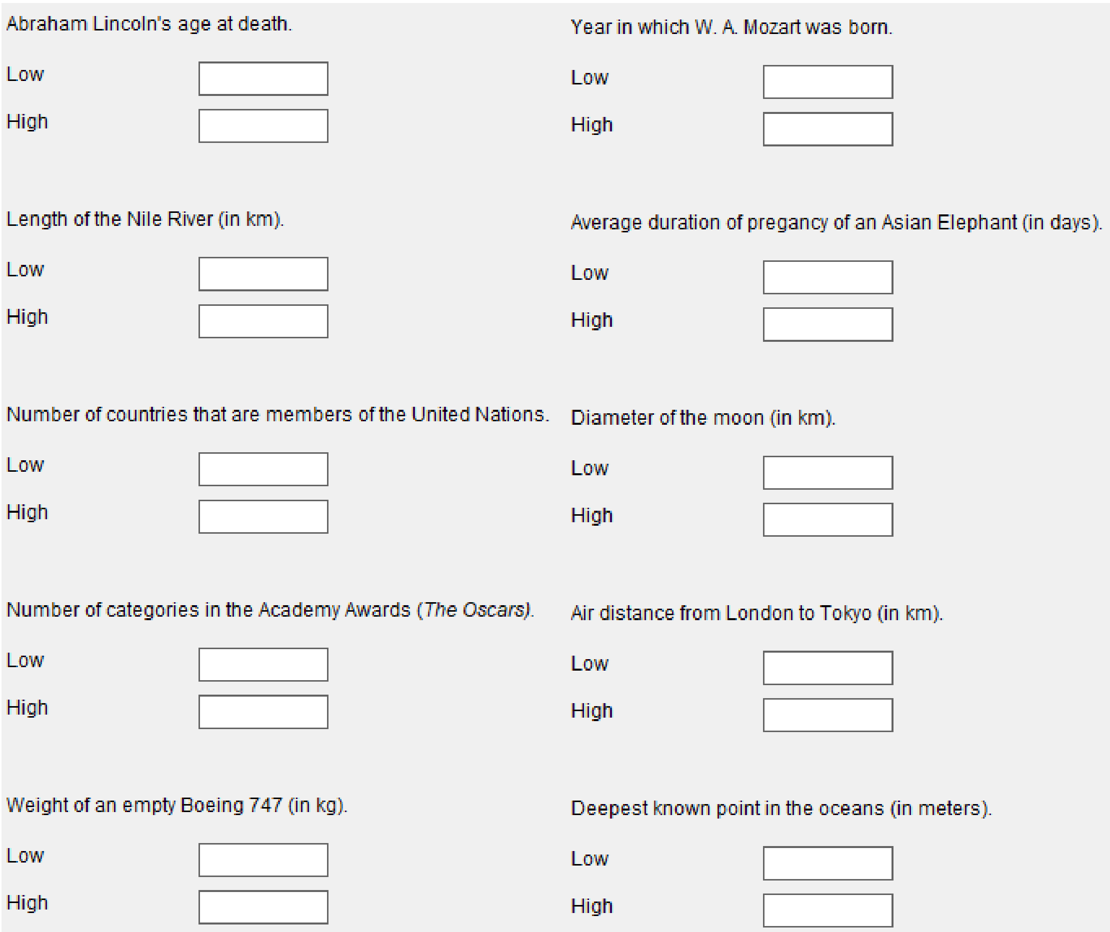

Figure C.14: Follow-up questionnaire, eliciting measure for self-esteem

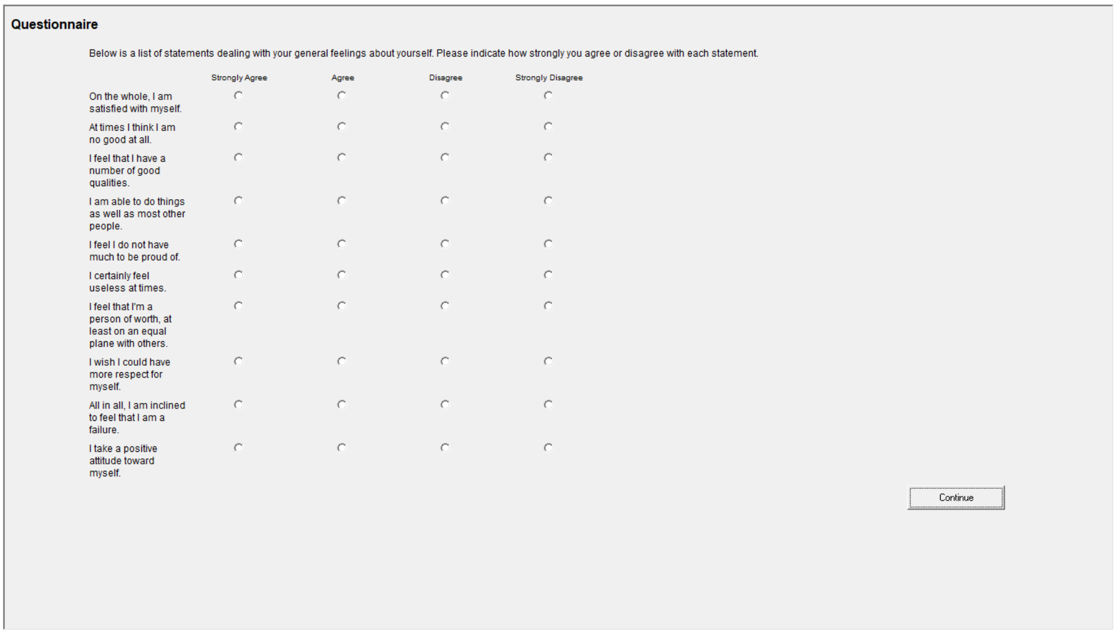

Questionnaire

Below is a list of statements dealing with your general feelings about yourself. Please indicate how strongly you agree or disagree with each statement.

	Strongly Agree	Agree	Disagree	Strongly Disagree
On the whole, I am satisfied with myself.	○	○	○	○
At times I think I am no good at all.	○	○	○	○
I feel that I have a number of good qualities.	○	○	○	○
I am able to do things as well as most other people.	○	○	○	○
I feel I do not have much to be proud of.	○	○	○	○
I certainly feel useless at times.	○	○	○	○
I feel that I'm a person of worth, at least on an equal plane with others.	○	○	○	○
I wish I could have more respect for myself.	○	○	○	○
All in all, I am inclined to feel that I am a failure.	○	○	○	○
I take a positive attitude toward myself.	○	○	○	○

Continue

Appendix D

Codes

D.1 Generating Data Bundles for the Experiment

```
#Import relevant packages
import os
import io
import random
import matplotlib.pyplot as plt
import matplotlib.patches as mpatches

#Number of Players can be adjusted as necessary
NumberofPeriods = 30
NumberofPlayers = 21
NumberofDatasets = NumberofPeriods*NumberofPlayers

#Defining some help-variables
ItemCounter = 0
AltItemCounter = 0
PlayerCounter = 0
PeriodCounter = 1

#Defining paths and filenames
dir = r"C:\Users\Moritz Mosenhauer\Academia Zeugs\PhD\Research\
    WhatIDontKnow\Data\Preparator2"
dir_input = dir + '\Input'
```

```
dir_output = dir + '\Output'
Inputs = 'CleanedReturnsMonthly.txt'
ResultName = "DataN" + str(NumberofPlayers) + "T" + str(
    NumberofPeriods) + "Monthly.txt"

#Reading in Stock Data
os.chdir(dir_input)
RawInputs = io.open(Inputs, 'r', encoding='utf-8').read().
    splitlines()
RawInputs.pop(0)
TotalInputLength = len(RawInputs)
os.chdir(dir_output)
with open(ResultName, 'w') as TargetData:
     TargetData.write("")

#Looping through stocks
for RawInput in RawInputs:
     ReturnsList = RawInput.split('\t')
     ReturnsList = ReturnsList[:-1]
     #Excluding stocks with less than 11 entries
     if len(ReturnsList) < 11:
          print ("LIST IS TOO SHORT, WE HAVE A PROBLEM!!!")
     #Determining total number of potential starting periods for
          hold-
     #and trade-stocks
     else:
          ItemCounter = int(ItemCounter) + len(ReturnsList) - 10
          AltItemCounter = int(AltItemCounter) + len(ReturnsList)

#Generate Databundle for each player and each period
for i in range(NumberofDatasets):
     PlotList = []
     PlayerCounter = int(PlayerCounter) + 1

     #Randomly picking hold-stock series of returns
```

```
StockIndex = random.randint(1,(ItemCounter))
#Randomly picking trade-stock series
AltStockIndex = random.randint(1,(AltItemCounter))


for RawInput in RawInputs:
    #Finding the relevant stock in list w.r.t. random pick
    ReturnsList = RawInput.split('\t')
    ReturnsList = ReturnsList[:-1]
    Checker = StockIndex - (len(ReturnsList) - 10)
    if Checker < 0:
        #Prepare plotting
        PlotList.append(ReturnsList[(StockIndex-1)])
        PlotList.append(ReturnsList[(StockIndex)])
        PlotList.append(ReturnsList[(StockIndex+1)])
        PlotList.append(ReturnsList[(StockIndex+2)])
        PlotList.append(ReturnsList[(StockIndex+3)])
        PlotList.append(ReturnsList[(StockIndex+4)])
        PlotList.append(ReturnsList[(StockIndex+5)])
        PlotList.append(ReturnsList[(StockIndex+6)])
        PlotList.append(ReturnsList[(StockIndex+7)])
        PlotList.append(ReturnsList[(StockIndex+8)])
        fig = plt.figure()
        ax = fig.add_subplot(111)

        #Set axes according to return-range
        if all( abs(float(entry))<30  for entry in PlotList
            ):
            ax.set_ylim([-31,31])
            yticklocat = [-30, -15, 0, 15, 30]
            ytickers = [r"-30%", r"-15%", r"0%", r"15%", r"
                30%"]
            ax.set_yticks(yticklocat)
            ax.set_yticklabels(ytickers, fontsize= 18)
        else:
```

```
    if all( abs(float(entry))<0.50  for entry in
        PlotList ):
         ax.set_ylim([-51,51])
         yticklocat = [-50, -25, 0, 25, 50]
         ytickers = [r"-50%", r"-25%", r"0%", r"25%"
             , r"50%"]
         ax.set_yticks(yticklocat)
         ax.set_yticklabels(ytickers, fontsize= 18)
    else:
         ax.set_ylim([-151,151])
         yticklocat = [-150, -75, 0, 75, 150]
         ytickers = [r"-150%", r"-75%", r"0%", r"75%
             ", r"150%"]
         ax.set_yticks(yticklocat)
         ax.set_yticklabels(ytickers, fontsize= 18)

#Adjust Layout
xticklocat = [0, 3, 6, 9]
xtickers = ["10 months \n prior", "7 months \n 
    prior", "4 months \n prior", r"last month"]
ax.set_xticks(xticklocat)
ax.set_xticklabels(xtickers, fontsize= 14)
plt.xticks(rotation=70)

Dashedx = [0, 9]
Dashedy = [0, 0]

label = ax.set_ylabel('Stock Returns', fontsize =
    17, rotation="horizontal")
ax.yaxis.set_label_coords(-0.025, 1.015)

plt.tight_layout()

#Plot and Save Figure
ax.plot(PlotList, color='b', linewidth='3')
```

```
ax.plot(Dashedx, Dashedy, 'k--', linewidth='2')

plt.savefig("StockHistN" + str(PlayerCounter) + "T"
    + str(PeriodCounter) + ".png", dpi=150)

#Determine unbroken series of recent positive or
    negative returns
RevStockReturns = reversed(PlotList)

PosCounter = 0
NegCounter = 0
TunnelCounter = 0
for entry in RevStockReturns:
    if float(entry) < 0:
        NegCounter = int(NegCounter) + 1
    else:
        break
RevStockReturns2 = reversed(PlotList)
for entry2 in RevStockReturns2:
    if float(entry2) >= 0:
        PosCounter = int(PosCounter) + 1
    else:
        break

if PosCounter > NegCounter:
    TunnelCounter = PosCounter
else:
    TunnelCounter = NegCounter

#Write all data into files
with open(ResultName, 'a') as TargetData:
    TargetData.write(str(ReturnsList[(StockIndex-1)
        ]))
    TargetData.write("; Return2 = ")
    TargetData.write(str(ReturnsList[(StockIndex)])
```

```
            )
        TargetData.write("; Return3 = ")
        TargetData.write(str(ReturnsList[(StockIndex+1)
            ]))
        TargetData.write("; Return4 = ")
        TargetData.write(str(ReturnsList[(StockIndex+2)
            ]))
        TargetData.write("; Return5 = ")
        TargetData.write(str(ReturnsList[(StockIndex+3)
            ]))
        TargetData.write("; Return6 = ")
        TargetData.write(str(ReturnsList[(StockIndex+4)
            ]))
        TargetData.write("; Return7 = ")
        TargetData.write(str(ReturnsList[(StockIndex+5)
            ]))
        TargetData.write("; Return8 = ")
        TargetData.write(str(ReturnsList[(StockIndex+6)
            ]))
        TargetData.write("; Return9 = ")
        TargetData.write(str(ReturnsList[(StockIndex+7)
            ]))
        TargetData.write("; Return10 = ")
        TargetData.write(str(ReturnsList[(StockIndex+8)
            ]))
        TargetData.write("; ReturnHold = ")
        TargetData.write(str(ReturnsList[(StockIndex+9)
            ]))
        TargetData.write("; ReturnSell = ")
    break
else:
    StockIndex = int(StockIndex) - (len(ReturnsList) -
        10)

#Find and write trade-stock return w.r.t. to random pick
```

```
    for RawInput in RawInputs:
        ReturnsList = RawInput.split('\t')
        ReturnsList = ReturnsList[:-1]
        AltChecker = AltStockIndex - len(ReturnsList)
        if AltChecker < 0:
            with open(ResultName, 'a') as TargetData:
                TargetData.write(str(ReturnsList[(AltStockIndex
                    -1)]))
            break
        else:
            AltStockIndex = int(AltStockIndex) - len(
                ReturnsList)

    #Advancing technical counters
    with open(ResultName, 'a') as TargetData:
        TargetData.write("; TunnelStimulus = ")
        TargetData.write(str(TunnelCounter))

    if PlayerCounter == int(NumberofPlayers):
        with open(ResultName, 'a') as TargetData:
            TargetData.write("\n")
        PlayerCounter = 0
        PeriodCounter = int(PeriodCounter) + 1
    else:
        with open(ResultName, 'a') as TargetData:
            TargetData.write("\t")
```

D.2 Formatting Data from Experiments for Analysis

D.2.1 Choice-Level Data

```
#Import Porgrammes
import os
import io
import math
```

```
import numpy as np

#Defining help-variables
SessionCounter = 0
SessionTicker = 0
IDCounter = 0
InfoRestrDummy = 0
ControlDummy = 0
FeedbackTreatDummy = 0
InfoRestrTreatDummy = 0
TreatInteractDummy = 0
TotalParticipants = 0
NBySessionList = []
ParticipantHelpCounter = 0

#Defining Strings - Paths and Core-Tableheaders
dir = r"C:\Users\Moritz Mosenhauer\Academia Zeugs\PhD\Research\
    WhatIDontKnow\Data\SessionData\SessionDataPreparator"
dir_input = dir + '\Input'
dir_output = dir + '\Output'
Inputs = 'CleanedReturnsMonthly.txt'
ResultName = "CollatedSessionData.txt"
CoreTableHeader = "EntryID\tSession\tPeriod\tSubject\tGroup\
    tProfit\tTotalProfit\tParticipate\tReturn1\tReturn2\tReturn3
    \tReturn4\tReturn5\tReturn6\tReturn7\tReturn8\tReturn9\
    tReturn10\tReturnHold\tReturnSell\tz\tInterPersz\
    tTunnelStimulus\tCommissionFee\tFakeCommssion\tTrading\
    tShowInfoMain\tShowInfoz\tN\tNumGroups\tMinPerGroupOne\
    tNumberPart\tGroupOne\tGroupTwo\tGroupThree\tGroupFour\
    tTimeContinueInstructionsControlOneOK\
    tTimeContinueInstructionsControlTwoOK\
    tTimeContinueInstructionsControlThreeOK\
    tTimeContinueInstructionsTreatmentIOneOK\
    tTimeContinueInstructionsTreatmentITwoOK\
    tTimeContinueInstructionsTreatmentIThreeOK\
```

```
    tTimeContinueInstructionsTreatmentIIOneOK\
    tTimeContinueInstructionsTreatmentIITwoOK\
    tTimeContinueInstructionsTreatmentIIThreeOK\
    tTimeContinueInstructionsInterIOneOK\
    tTimeContinueInstructionsInterITwoOK\
    tTimeContinueInstructionsInterIThreeOK\tPosPeriods\
    tNegPeriods\tTimeShowInstructionsMainControlOK\
    tTimeContinueMainControlOK\
    tTimeReturnToDecisionMainControlOK\
    tTimeShowInstructionsMainTreatIOK\tTimeContinueMainTreatIOK\
    tTimeReturnToDecisionMainTreatIOK\
    tTimeShowInstructionsMainTreatIIOK\
    tTimeContinueMainTreatIIOK\
    tTimeReturnToDecisionMainTreatIIOK\
    tTimeShowInstructionsMainInterIOK\tTimeContinueMainInterIOK\
    tTimeReturnToDecisionMainInterIOK\tAltProfit\
    tTimeContinueShowProfitsTreatIOK\
    tTimeContinueShowProfitsInterIOK\tInfoRestrDummy\
    tControlDummy\tFeedbackTreatDummy\tInfoRestrTreatDummy\
    tTreatInteractDummy\tGrossProfit\tNormProfit\tRootNormProfit
    \tLogNormProfit\tHistReturnsVariance\tHistReturnsStd"
CoreQuestHeader = "\tQuestSubject\tclient\tgender\tage\theight\
    tweight\twealth\tstocksOne\tstocksTwo\tspeeding\tgambling\
    tdyscalculia\triskOne\triskTwo\triskThree\triskFour\
    triskFive\triskSix\triskSeven\triskEight\triskNine\triskTen\
    tconfLowOne\tconfHighOne\tconfLowTwo\tconfHighTwo\
    tconfLowThree\tconfHighThree\tconfLowFour\tconfHighFour\
    tconfLowFive\tconfHighFive\tconfLowSix\tconfHighSix\
    tconfLowSeven\tconfHighSeven\tconfLowEight\tconfHighEight\
    tconfLowNine\tconfHighNine\tconfLowTen\tconfHighTen\
    tSelfEsteemOne\tSelfEsteemTwo\tSelfEsteemThree\
    tSelfEsteemFour\tSelfEsteemFive\tSelfEsteemSix\
    tSelfEsteemSeven\tSelfEsteemEight\tSelfEsteemNine\
    tSelfEsteemTen"
ExtraQuestHeader = "\tBinGenDummy\tMaleDummy\tFemaleDummy\
```

```
    tOthGenDummy\tBMI\tBMIDummyStrict\tBMIDummyExtreme\
    tStockExpDummyOne\tStockExpDummyTwo\tGamblingIndex\
    tDyscalculiaDummy\tRiskSeekIndex\tRiskConsistDummy\
    tOverconfIndex\tHighOverconfDummy\tLowSelfEsteemIndex\
    tFeedTreatMaleDummyInter\tInforestTreatMaleDummyInter\
    tTreatInterMaleDummyInter\tFeedTreatAgeInter\
    tInforestTreatAgeInter\tTreatInterAgeInter\
    tFeedTreatWealthInter\tInforestTreatWealthInter\
    tTreatInterWealthInter\tFeedTreatStockExpOneInter\
    tInforestTreatStockExpOneInter\tTreatInterStockExpOneInter\
    tFeedTreatStockExpTwoInter\tInforestTreatStockExpTwoInter\
    tTreatInterStockExpTwoInter\tFeedTreatSpeedingInter\
    tInforestTreatSpeedingInter\tTreatInterSpeedingInter\
    tFeedTreatGamblIndexInter\tInforestTreatGamblIndexInter\
    tTreatInterGamblIndexInter\tFeedTreatDyscalculiaInter\
    tInforestTreatDyscalculiaInter\tTreatInterDyscalculiaInter\
    tFeedTreatRiskSeekIndexInter\
    tInforestTreatRiskSeekIndexInter\
    tTreatInterRiskSeekIndexInter\tFeedTreatConfIndexInter\
    tInforestTreatConfIndexInter\tTreatInterConfIndexInter\
    tFeedTreatConfDummyInter\tInforestTreatConfDummyInter\
    tTreatInterConfDummyInter\tFeedTreatLowSelfEstIndexInter\
    tInforestTreatLowSelfEstIndexInter\
    tTreatInterLowSelfEstIndexInter\tFakeEntryDummy"

#Make for Browse through all datasets to determine number of
    participants (total and per session)
for Prefilename in os.listdir(dir_input):
    PartCompareList = []
    os.chdir(dir_input)
    if str(Prefilename).endswith('.xls'):
        SessionCounter = SessionCounter + 1
        PreRawInputs = io.open(Prefilename, 'r', encoding='utf
            -8').read().splitlines()
        for PreRawInput in PreRawInputs:
```

```
            PreDataEntryList = PreRawInput.split('\t')
            if PreDataEntryList[3] == "Subject":
                pass
            else:
                if PreDataEntryList[2] == "session":
                    PartCompareList.append(int(PreDataEntryList
                        [3]))
        NBySessionList.append(max(PartCompareList))
        TotalParticipants = TotalParticipants + max(
            PartCompareList)

print ("The total number of participants is " + str(
    TotalParticipants))

#Table-Headers for Session-, Participants- and Time-Dummies
SessionHeader = ""
for SessionIndex in range(SessionCounter):
    SessionHeader = SessionHeader + "\tSession" + str((
        SessionIndex+1))

ParticipantHeader = ""
for ParticipantIndex in range(TotalParticipants):
    ParticipantHeader = ParticipantHeader + "\tParticip" + str
        ((ParticipantIndex+1))

PeriodHeader = ""
for PeriodHeaderIndex in range(30):
    PeriodHeader = PeriodHeader + "\tPeriod" + str((
        PeriodHeaderIndex+1))
FinalTableHeader = CoreTableHeader + SessionHeader +
    PeriodHeader + ParticipantHeader + CoreQuestHeader +
    ExtraQuestHeader + "\n"

#Writing Tableheader
os.chdir(dir_output)
```

```python
with open(ResultName, 'w') as TargetData:
    TargetData.write(FinalTableHeader)

#Start Main-Loop
for filename in os.listdir(dir_input):
    os.chdir(dir_input)
    #Open Tables
    if str(filename).endswith('.xls'):
        SessionTicker = SessionTicker + 1
        #Break into Entrylines
        print(filename)
        RawInputs = io.open(filename, 'r', encoding='utf-8').
            read().splitlines()
        for RawInput in RawInputs:
            #Break Lines into Datapoints
            DataEntryList = RawInput.split('\t')
            #Only treat relevant lines in original table
            if DataEntryList[3] == "Period":
                pass
            else:
                if DataEntryList[2] == "subjects":
                    HistoricReturnsList = []

                    IDCounter = IDCounter + 1

                    #Creating Fixed Effects Dummies
                    SessionDummyList = []
                    for i in range(SessionCounter-1):
                        SessionDummyList.append("0\t")
                    SessionDummyList.insert( (SessionTicker-1),
                        "1\t")
                    SessionDummyEntry = ''.join(
                        SessionDummyList)

                    TimeDummyList = []
```

```
for i in range(30-1):
    TimeDummyList.append("0\t")
TimeDummyList.insert( (int(DataEntryList
   [3])-1), "1\t")
TimeDummyEntry = ''.join(TimeDummyList)

ParticipantDummyList = []
for i in range(TotalParticipants-1):
    ParticipantDummyList.append("0\t")
ParticipantDummyList.insert( (
   ParticipantHelpCounter + int(
   DataEntryList[4])-1), "1\t")
ParticipantDummyEntry = ''.join(
   ParticipantDummyList)

#Creating Treatment Dummies
ControlDummy = 0
FeedbackTreatDummy = 0
InfoRestrTreatDummy = 0
TreatInteractDummy = 0

if int(DataEntryList[5]) == 1:
    ControlDummy = 1
if (int(DataEntryList[5]) == 2) or (int(
   DataEntryList[5]) == 4):
    FeedbackTreatDummy = 1
if (int(DataEntryList[5]) == 3) or (int(
   DataEntryList[5]) == 4):
    InfoRestrTreatDummy = 1
if int(DataEntryList[5]) == 4:
    TreatInteractDummy = 1

#Profit Manipulations
if DataEntryList[26]=="0":
    GrossProfit = DataEntryList[6]
```

```
else:
    GrossProfit = DataEntryList[20]

MinProfit = -79.9748
MaxProfit = 171.7857
NormProfit = (float(DataEntryList[6])-
   MinProfit)/(MaxProfit-MinProfit)
RootNormProfit = math.sqrt(NormProfit)
if not NormProfit == 0:
    LogNormProfit = math.log(NormProfit)
else:
    LogNormProfit = ""

for i in range(10):
    HistoricReturnsList.append(float(
       DataEntryList[9 + i]))
HistReturnsVariance = np.var(
   HistoricReturnsList)
HistReturnsStd = np.std(HistoricReturnsList
   )

#Creating Information-Restriction Dummy
if int(DataEntryList[23]) > int(
   DataEntryList[21]):
    InfoRestrDummy = 0
else:
    InfoRestrDummy = 1

DataEntryList = DataEntryList[:-1]
DataEntryListChopped = DataEntryList[3:]

os.chdir(dir_output)
#Writing MainData
with open(ResultName, 'a') as TargetData:
    TargetData.write(str(IDCounter))
```

```
TargetData.write("\t")
TargetData.write(str(SessionTicker))
TargetData.write("\t")
TargetData.write("\t".join(
    DataEntryList[3:]))
TargetData.write("\t")
TargetData.write(str(InfoRestrDummy))
TargetData.write("\t")
TargetData.write(str(ControlDummy))
TargetData.write("\t")
TargetData.write(str(FeedbackTreatDummy
    ))
TargetData.write("\t")
TargetData.write(str(
    InfoRestrTreatDummy))
TargetData.write("\t")
TargetData.write(str(TreatInteractDummy
    ))
TargetData.write("\t")
TargetData.write(str(GrossProfit))
TargetData.write("\t")
TargetData.write(str(NormProfit))
TargetData.write("\t")
TargetData.write(str(RootNormProfit))
TargetData.write("\t")
TargetData.write(str(LogNormProfit))
TargetData.write("\t")
TargetData.write(str(
    HistReturnsVariance))
TargetData.write("\t")
TargetData.write(str(HistReturnsStd))
TargetData.write("\t")
TargetData.write(SessionDummyEntry)
TargetData.write(TimeDummyEntry)
TargetData.write(ParticipantDummyEntry)
```

```
os.chdir(dir_input)
#Loop for finding relevant questionnaire-
    data for each main-entry
for questfilename in os.listdir(dir_input):
    #Open and split into data-points
    if (str(questfilename).startswith(str(
        filename[:11])) and str(
        questfilename).endswith('.txt')):
        QuestInputs = io.open(questfilename
            , 'r', encoding='utf-8').read().
            splitlines()
        ParticQuestEntry = QuestInputs[int(
            DataEntryList[4])]
        QuestDataEntries = ParticQuestEntry
            .split('\t')

        #Creating two types of gender-
            dummies
        if QuestDataEntries[2] == "male":
            BinaryGenderDummy = "1"
            MaleDummy = 1
            FemaleDummy = "0"
            OtherGenderDummy = "0"
        else:
            if QuestDataEntries[2] == "
                female":
                BinaryGenderDummy = "0"
                MaleDummy = 0
                FemaleDummy = "1"
                OtherGenderDummy = "0"
            else:
                BinaryGenderDummy = "-"
                MaleDummy = 0
                FemaleDummy = "0"
```

```
            OtherGenderDummy = "1"

    #Calculating BMI and two
        interpretative dummies
    if not (QuestDataEntries[4]=="-" or
         QuestDataEntries[4]=="0" or
        QuestDataEntries[5]=="-" or
        QuestDataEntries[5]=="0"):
        BMI = 10000*int(
            QuestDataEntries[5])/float(
            math.pow(int(
            QuestDataEntries[4]), 2))
        if ( 18.5< BMI and 25 > BMI ):
            BMIDummyStrict = "0"
        else:
            BMIDummyStrict = "1"
        if ( 17< BMI and 30 > BMI ):
            BMIDummyExtreme = "0"
        else:
            BMIDummyExtreme = "1"
    else:
        BMI = ""
        BMIDummyStrict = ""
        BMIDummyExtreme = ""

    #Stock-Trading Experience Dummies
    if QuestDataEntries[7] == "yes":
        StockExpDummyOne = 1
    else:
        StockExpDummyOne = 0
    if QuestDataEntries[8] == "yes":
        StockExpDummyTwo = 1
    else:
        StockExpDummyTwo = 0
```

```
#Gambling Dummies
if QuestDataEntries[10] == "never":
    GamblingIndex = 0
if QuestDataEntries[10] == "1-5 
   times":
    GamblingIndex = 1/float(3)
if QuestDataEntries[10] == "6 - 10 
   times":
    GamblingIndex = 2/float(3)
if QuestDataEntries[10] == "more 
   than 10 times":
    GamblingIndex = 1

#Dyscalculia-Dummies
if QuestDataEntries[11] == "yes":
    DyscalculiaDummy = 1
else:
    DyscalculiaDummy = 0

#Creating Risk-Seeking Index and
   Dummy whether
#statements are consistent with
   rationality
RiskSeekIndex = 0
previouselement = None
ConsistencyHitCounter = 0
RiskConsistDummy = "1"
for element in QuestDataEntries
   [12:22]:
    RiskSeekIndex = RiskSeekIndex +
        int(element)
    if not previouselement == None:
        if not previouselement ==
           element:
            ConsistencyHitCounter =
```

```
                    ConsistencyHitCounter
                    + 1
        previouselement = element
    if ConsistencyHitCounter > 1:
        RiskConsistDummy = "0"
    RiskSeekIndex = RiskSeekIndex/float
        (10)

    #Creating Overconfidence-Index
    ConfHitCounter = 0
    ConfProblemDummy = 0
    HighOverconfDummy = 0
    if not (QuestDataEntries[22]=="-"
        or QuestDataEntries[23]=="-"):
        if not int(QuestDataEntries
            [22]) <= 56 <= int(
            QuestDataEntries[23]):
            ConfHitCounter =
                ConfHitCounter + 1
    else:
        ConfProblemDummy = 1
    if not (QuestDataEntries[24]=="-"
        or QuestDataEntries[25]=="-"):
        if not int(QuestDataEntries
            [24]) <= 6853 <= int(
            QuestDataEntries[25]):
            ConfHitCounter =
                ConfHitCounter + 1
    else:
        ConfProblemDummy = 1
    if not (QuestDataEntries[26]=="-"
        or QuestDataEntries[27]=="-"):
        if not int(QuestDataEntries
            [26]) <= 193 <= int(
```

```
        QuestDataEntries[27]):
         ConfHitCounter =
            ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[28]=="-"
   or QuestDataEntries[29]=="-"):
    if not int(QuestDataEntries
        [28]) <= 24 <= int(
        QuestDataEntries[29]):
         ConfHitCounter =
            ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[30]=="-"
   or QuestDataEntries[31]=="-"):
    if not int(QuestDataEntries
        [30]) <= 380000 <= int(
        QuestDataEntries[31]):
         ConfHitCounter =
            ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[32]=="-"
   or QuestDataEntries[33]=="-"):
    if not int(QuestDataEntries
        [32]) <= 1756 <= int(
        QuestDataEntries[33]):
         ConfHitCounter =
            ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[34]=="-"
   or QuestDataEntries[35]=="-"):
    if not int(QuestDataEntries
```

```
            [34]) <= 6000 <= int(
            QuestDataEntries[36]):
             ConfHitCounter =
                 ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[36]=="-"
    or QuestDataEntries[37]=="-"):
     if not int(QuestDataEntries
            [36]) <= 3474 <= int(
            QuestDataEntries[37]):
             ConfHitCounter =
                 ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[38]=="-"
    or QuestDataEntries[39]=="-"):
     if not int(QuestDataEntries
            [38]) <= 9562 <= int(
            QuestDataEntries[39]):
             ConfHitCounter =
                 ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[40]=="-"
    or QuestDataEntries[41]=="-"):
     if not int(QuestDataEntries
            [40]) <= 10994 <= int(
            QuestDataEntries[41]):
             ConfHitCounter =
                 ConfHitCounter + 1
else:
     ConfProblemDummy = 1


if ConfProblemDummy == 0:
```

```
    OverconfIndex = ConfHitCounter/
        float(10)
    if OverconfIndex > 0.7:
        HighOverconfDummy = 1
else:
    HighOverconfDummy = ""
    OverconfIndex = ""

#Creating Low-SelfesteemIndex
SelfEsteemSumOne = 15 - int(
    QuestDataEntries[42]) - int(
    QuestDataEntries[44]) - int(
    QuestDataEntries[45]) - int(
    QuestDataEntries[48]) - int(
    QuestDataEntries[51])
SelfEsteemSumTwo = int(
    QuestDataEntries[43]) + int(
    QuestDataEntries[46]) + int(
    QuestDataEntries[47]) + int(
    QuestDataEntries[49]) + int(
    QuestDataEntries[50])
LowSelfEsteemIndex = (
    SelfEsteemSumOne +
    SelfEsteemSumTwo)/float(30)



### Creating Personal/Treatment-
    Interactions ###

#Male Dummy
FeedTreatMaleDummyInter =
    FeedbackTreatDummy*MaleDummy
InforestTreatMaleDummyInter =
    InfoRestrTreatDummy*MaleDummy
TreatInterMaleDummyInter =
```

```
    TreatInteractDummy*MaleDummy

#Age
Age = float(QuestDataEntries[3])
FeedTreatAgeInter =
    FeedbackTreatDummy*Age
InforestTreatAgeInter =
    InfoRestrTreatDummy*Age
TreatInterAgeInter =
    TreatInteractDummy*Age

#Wealth
Wealth = float(QuestDataEntries[6])
FeedTreatWealthInter =
    FeedbackTreatDummy*Wealth
InforestTreatWealthInter =
    InfoRestrTreatDummy*Wealth
TreatInterWealthInter =
    TreatInteractDummy*Wealth

#Stockexperience One
FeedTreatStockExpOneInter =
    FeedbackTreatDummy*
    StockExpDummyOne
InforestTreatStockExpOneInter =
    InfoRestrTreatDummy*
    StockExpDummyOne
TreatInterStockExpOneInter =
    TreatInteractDummy*
    StockExpDummyOne

#Stockexperience Two
FeedTreatStockExpTwoInter =
    FeedbackTreatDummy*
    StockExpDummyTwo
```

```
InforestTreatStockExpTwoInter =
    InfoRestrTreatDummy*
    StockExpDummyTwo
TreatInterStockExpTwoInter =
    TreatInteractDummy*
    StockExpDummyTwo

#Speeding
Speeding = float(QuestDataEntries
    [9])
FeedTreatSpeedingInter =
    FeedbackTreatDummy*Speeding
InforestTreatSpeedingInter =
    InfoRestrTreatDummy*Speeding
TreatInterSpeedingInter =
    TreatInteractDummy*Speeding

#GamblingIndex
FeedTreatGamblIndexInter =
    FeedbackTreatDummy*GamblingIndex
InforestTreatGamblIndexInter =
    InfoRestrTreatDummy*
    GamblingIndex
TreatInterGamblIndexInter =
    TreatInteractDummy*GamblingIndex

#Dyscalculia
FeedTreatDyscalculiaInter =
    FeedbackTreatDummy*
    DyscalculiaDummy
InforestTreatDyscalculiaInter =
    InfoRestrTreatDummy*
    DyscalculiaDummy
TreatInterDyscalculiaInter =
    TreatInteractDummy*
```

```
    DyscalculiaDummy

#RiskSeekIndex
FeedTreatRiskSeekIndexInter =
    FeedbackTreatDummy*RiskSeekIndex
InforestTreatRiskSeekIndexInter =
    InfoRestrTreatDummy*
    RiskSeekIndex
TreatInterRiskSeekIndexInter =
    TreatInteractDummy*RiskSeekIndex

#Overconfidence Index
if not OverconfIndex == "":
     FeedTreatConfIndexInter =
         FeedbackTreatDummy*
         OverconfIndex
     InforestTreatConfIndexInter =
         InfoRestrTreatDummy*
         OverconfIndex
     TreatInterConfIndexInter =
         TreatInteractDummy*
         OverconfIndex
else:
     FeedTreatConfInter = ""
     InforestTreatConfInter = ""
     TreatInterConfInter = ""

#Overconfidence Dummy
if not HighOverconfDummy == "":
     FeedTreatConfDummyInter =
         FeedbackTreatDummy*
         HighOverconfDummy
     InforestTreatConfDummyInter =
         InfoRestrTreatDummy*
         HighOverconfDummy
```

```
        TreatInterConfDummyInter =
            TreatInteractDummy*
            HighOverconfDummy
    else:
        FeedTreatConfDummyInter = ""
        InforestTreatConfDummyInter = "
            "
        TreatInterConfDummyInter = ""

    #LowSelfEsteemIndex
    FeedTreatLowSelfEstIndexInter =
        FeedbackTreatDummy*
        LowSelfEsteemIndex
    InforestTreatLowSelfEstIndexInter =
        InfoRestrTreatDummy*
        LowSelfEsteemIndex
    TreatInterLowSelfEstIndexInter =
        TreatInteractDummy*
        LowSelfEsteemIndex

    #Creating Fake-Entry Dummy
    FakeEntryDummy = 0
    #if (str(OverconfIndex)=="1.0" and
        QuestDataEntries[32]=="0" and
        BMI=="-"):
    if ((str(DataEntryList[0])=="171012
        _1723" and str(DataEntryList[4])
        =="10") or (str(DataEntryList
        [0])=="171115_1730" and str(
        DataEntryList[4])=="13") or (str
        (DataEntryList[0])=="171018_1718
        " and str(DataEntryList[4])=="28
        ")):
        FakeEntryDummy = 1
        #print("We got Subject " + str(
```

```
                    DataEntryList[4]) + "!")

#Writing everything down
os.chdir(dir_output)
with open(ResultName, 'a') as TargetData:
    TargetData.write(ParticQuestEntry)

with open(ResultName, 'a') as TargetData:
    TargetData.write("\t")
    TargetData.write(BinaryGenderDummy)
    TargetData.write("\t")
    TargetData.write(str(MaleDummy))
    TargetData.write("\t")
    TargetData.write(FemaleDummy)
    TargetData.write("\t")
    TargetData.write(OtherGenderDummy)
    TargetData.write("\t")
    TargetData.write(str(BMI))
    TargetData.write("\t")
    TargetData.write(BMIDummyStrict)
    TargetData.write("\t")
    TargetData.write(BMIDummyExtreme)
    TargetData.write("\t")
    TargetData.write(str(StockExpDummyOne))
    TargetData.write("\t")
    TargetData.write(str(StockExpDummyTwo))
    TargetData.write("\t")
    TargetData.write(str(GamblingIndex))
    TargetData.write("\t")
    TargetData.write(str(DyscalculiaDummy))
    TargetData.write("\t")
    TargetData.write(str(RiskSeekIndex))
    TargetData.write("\t")
    TargetData.write(RiskConsistDummy)
    TargetData.write("\t")
```

```
TargetData.write(str(OverconfIndex))
TargetData.write("\t")
TargetData.write(str(HighOverconfDummy)
    )
TargetData.write("\t")
TargetData.write(str(LowSelfEsteemIndex
    ))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(FeedTreatAgeInter)
    )
TargetData.write("\t")
TargetData.write(str(
    InforestTreatAgeInter))
TargetData.write("\t")
TargetData.write(str(TreatInterAgeInter
    ))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatWealthInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatWealthInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterWealthInter))
TargetData.write("\t")
```

```
TargetData.write(str(
    FeedTreatStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatSpeedingInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatSpeedingInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterSpeedingInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatGamblIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatGamblIndexInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterGamblIndexInter))
```

```
TargetData.write("\t")
TargetData.write(str(
    FeedTreatDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatConfIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatConfIndexInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterConfIndexInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatConfDummyInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatConfDummyInter))
TargetData.write("\t")
TargetData.write(str(
```

```
                TreatInterConfDummyInter))
            TargetData.write("\t")
            TargetData.write(str(
                FeedTreatLowSelfEstIndexInter))
            TargetData.write("\t")
            TargetData.write(str(
                InforestTreatLowSelfEstIndexInter))
            TargetData.write("\t")
            TargetData.write(str(
                TreatInterLowSelfEstIndexInter))
            TargetData.write("\t")
            TargetData.write(str(FakeEntryDummy))

        #Next Line in Data
        with open(ResultName, 'a') as TargetData:
            TargetData.write("\n")
    ParticipantHelpCounter = ParticipantHelpCounter +
        NBySessionList[SessionTicker-1]
```

D.2.2 Individual-Level Data

```
#Import Porgrammes
import os
import io
import math
import numpy as np

#Defining help-variables
SessionCounter = 0
SessionTicker = 0
IDCounter = 0
InfoRestrDummy = 0
ControlDummy = 0
FeedbackTreatDummy = 0
InfoRestrTreatDummy = 0
```

```
TreatInteractDummy = 0
TotalParticipants = 0
NBySessionList = []
ParticipantHelpCounter = 0

def Average(lst):
    return sum(lst) / len(lst)

#Defining Strings - Paths and Core-Tableheaders
dir = r"C:\Users\Moritz Mosenhauer\Academia Zeugs\PhD\Research\
    WhatIDontKnow\Data\SessionData\
    SessionDataPreparatorAggregated"
dir_input = dir + '\Input'
dir_output = dir + '\Output'
Inputs = 'CleanedReturnsMonthly.txt'
ResultName = "CollatedSessionDataAggregated.txt"
CoreTableHeader = "EntryID\tSession\tRandomGroup\tControlDummy\
    tFeedbackTreatDummy\tInfoRestrTreatDummy\tTreatInteractDummy
    \tAverageNetReturns\tAverageTrading\tAverageGrossReturns\
    tNetReturnVariance\tNetReturnStd\tGrossReturnVariance\
    tGrossReturnStd\tTradingQuintile\tFakeDummy"
CoreQuestHeader = "\tQuestSubject\tclient\tgender\tage\theight\
    tweight\twealth\tstocksOne\tstocksTwo\tspeeding\tgambling\
    tdyscalculia\triskOne\triskTwo\triskThree\triskFour\
    triskFive\triskSix\triskSeven\triskEight\triskNine\triskTen\
    tconfLowOne\tconfHighOne\tconfLowTwo\tconfHighTwo\
    tconfLowThree\tconfHighThree\tconfLowFour\tconfHighFour\
    tconfLowFive\tconfHighFive\tconfLowSix\tconfHighSix\
    tconfLowSeven\tconfHighSeven\tconfLowEight\tconfHighEight\
    tconfLowNine\tconfHighNine\tconfLowTen\tconfHighTen\
    tSelfEsteemOne\tSelfEsteemTwo\tSelfEsteemThree\
    tSelfEsteemFour\tSelfEsteemFive\tSelfEsteemSix\
    tSelfEsteemSeven\tSelfEsteemEight\tSelfEsteemNine\
    tSelfEsteemTen"
ExtraQuestHeader = "\tBinGenDummy\tMaleDummy\tFemaleDummy\
```

```
    tOthGenDummy\tBMI\tBMIDummyStrict\tBMIDummyExtreme\
    tStockExpDummyOne\tStockExpDummyTwo\tGamblingIndex\
    tDyscalculiaDummy\tRiskSeekIndex\tRiskConsistDummy\
    tOverconfIndex\tHighOverconfDummy\tLowSelfEsteemIndex\
    tFeedTreatMaleDummyInter\tInforestTreatMaleDummyInter\
    tTreatInterMaleDummyInter\tFeedTreatAgeInter\
    tInforestTreatAgeInter\tTreatInterAgeInter\
    tFeedTreatWealthInter\tInforestTreatWealthInter\
    tTreatInterWealthInter\tFeedTreatStockExpOneInter\
    tInforestTreatStockExpOneInter\tTreatInterStockExpOneInter\
    tFeedTreatStockExpTwoInter\tInforestTreatStockExpTwoInter\
    tTreatInterStockExpTwoInter\tFeedTreatSpeedingInter\
    tInforestTreatSpeedingInter\tTreatInterSpeedingInter\
    tFeedTreatGamblIndexInter\tInforestTreatGamblIndexInter\
    tTreatInterGamblIndexInter\tFeedTreatDyscalculiaInter\
    tInforestTreatDyscalculiaInter\tTreatInterDyscalculiaInter\
    tFeedTreatRiskSeekIndexInter\
    tInforestTreatRiskSeekIndexInter\
    tTreatInterRiskSeekIndexInter\tFeedTreatConfIndexInter\
    tInforestTreatConfIndexInter\tTreatInterConfIndexInter\
    tFeedTreatConfDummyInter\tInforestTreatConfDummyInter\
    tTreatInterConfDummyInter\tFeedTreatLowSelfEstIndexInter\
    tInforestTreatLowSelfEstIndexInter\
    tTreatInterLowSelfEstIndexInter\tFakeEntryDummy"

#Browse through all datasets to determine number of
    participants (total and per session)
for Prefilename in os.listdir(dir_input):
    PartCompareList = []
    os.chdir(dir_input)
    if str(Prefilename).endswith('.xls'):
        SessionCounter = SessionCounter + 1
        PreRawInputs = io.open(Prefilename, 'r', encoding='utf
            -8').read().splitlines()
        for PreRawInput in PreRawInputs:
```

```
            PreDataEntryList = PreRawInput.split('\t')
            if PreDataEntryList[3] == "Subject":
                pass
            else:
                if PreDataEntryList[2] == "session":
                    PartCompareList.append(int(PreDataEntryList
                        [3]))
        NBySessionList.append(max(PartCompareList))
        TotalParticipants = TotalParticipants + max(
            PartCompareList)

print ("The total number of participants is " + str(
    TotalParticipants))

#Table-Headers for Session-, Participants- and Time-Dummies
SessionHeader = ""
for SessionIndex in range(SessionCounter):
    SessionHeader = SessionHeader + "\tSession" + str((
        SessionIndex+1))

FinalTableHeader = CoreTableHeader + SessionHeader +
    CoreQuestHeader + ExtraQuestHeader + "\n"

#Writing Tableheader
os.chdir(dir_output)
with open(ResultName, 'w') as TargetData:
    TargetData.write(FinalTableHeader)

#Start Main-Loop
for filename in os.listdir(dir_input):
    os.chdir(dir_input)
    #Open Tables
    if str(filename).endswith('.xls'):
        SessionTicker = SessionTicker + 1
        #Break into Entrylines
```

```
print(filename)
RawInputs = io.open(filename, 'r', encoding='utf-8').
    read().splitlines()
for Participant in range(NBySessionList[SessionTicker
    -1]):
     #Creating Help Variables
     ParticipProfitList = []
     ParticipTradingList = []
     ParticipGrossProfitList = []
     FakeDummy = 0
     TreatmentHitDummy = 0

     #Creating Treatment Dummies
     ControlDummy = 0
     FeedbackTreatDummy = 0
     InfoRestrTreatDummy = 0
     TreatInteractDummy = 0

     Participant = str(int(Participant) + 1)
     for RawInput in RawInputs:
          #Break Lines into Datapoints
          DataEntryList = RawInput.split('\t')
          #Only treat relevant lines in original table
          if DataEntryList[3] == "Period":
               pass
          else:
               if DataEntryList[2] == "subjects":
                    if DataEntryList[4] == Participant:
                         ParticipProfitList.append(float(
                             DataEntryList[6]))
                         ParticipTradingList.append(float(
                             DataEntryList[26]))
                         #Profit Manipulations
                         if DataEntryList[26]=="0":
                              GrossProfit = DataEntryList[6]
```

```
else:
    GrossProfit = DataEntryList[20]
ParticipGrossProfitList.append(
    float(GrossProfit))

#Determine Treat - Only once per
    Participant
if TreatmentHitDummy == 0:
    if int(DataEntryList[5]) == 1:
        ControlDummy = 1
    if (int(DataEntryList[5]) == 2)
        or (int(DataEntryList[5])
        == 4):
        FeedbackTreatDummy = 1
    if (int(DataEntryList[5]) == 3)
        or (int(DataEntryList[5])
        == 4):
        InfoRestrTreatDummy = 1
    if int(DataEntryList[5]) == 4:
        TreatInteractDummy = 1

    #Identifying Synthetic-Entries
    if DataEntryList[0] == "171012
        _1723" and DataEntryList[4]
        == "10":
        FakeDummy = 1
    if DataEntryList[0] == "171115
        _1730" and DataEntryList[4]
        == "13":
        FakeDummy = 1
    if DataEntryList[0] == "171018
        _1718" and DataEntryList[4]
        == "28":
        FakeDummy = 1
```

```
                RandomGroup = DataEntryList[5]

                TreatmentHitDummy = 1

#Outcome Variables
AverageProfit = Average(ParticipProfitList)
AverageTrading = Average(ParticipTradingList)
AverageGrossProfit = Average(
   ParticipGrossProfitList)

NetReturnVariance = np.var(ParticipProfitList)
NetReturnStd = np.std(ParticipProfitList)
GrossReturnVariance = np.var(
   ParticipGrossProfitList)
GrossReturnStd = np.std(ParticipGrossProfitList)

TradingQuintile = 5
if AverageTrading < 0.5:
    TradingQuintile = 4
if AverageTrading < 0.366667:
    TradingQuintile = 3
if AverageTrading < 0.266667:
    TradingQuintile = 2
if AverageTrading < 0.166667:
    TradingQuintile = 1

IDCounter = IDCounter + 1

#Creating Dummy-Data
SessionDummyList = []
for i in range(SessionCounter-1):
    SessionDummyList.append("0\t")
SessionDummyList.insert( (SessionTicker-1), "1\t")
SessionDummyEntry = ''.join(SessionDummyList)
```

```
os.chdir(dir_output)
#Writing MainData
with open(ResultName, 'a') as TargetData:
    TargetData.write(str(IDCounter))
    TargetData.write("\t")
    TargetData.write(str(SessionTicker))
    TargetData.write("\t")
    TargetData.write(str(RandomGroup))
    TargetData.write("\t")
    TargetData.write(str(ControlDummy))
    TargetData.write("\t")
    TargetData.write(str(FeedbackTreatDummy))
    TargetData.write("\t")
    TargetData.write(str(InfoRestrTreatDummy))
    TargetData.write("\t")
    TargetData.write(str(TreatInteractDummy))
    TargetData.write("\t")
    TargetData.write(str(AverageProfit))
    TargetData.write("\t")
    TargetData.write(str(AverageTrading))
    TargetData.write("\t")
    TargetData.write(str(AverageGrossProfit))
    TargetData.write("\t")
    TargetData.write(str(NetReturnVariance))
    TargetData.write("\t")
    TargetData.write(str(NetReturnStd))
    TargetData.write("\t")
    TargetData.write(str(GrossReturnVariance))
    TargetData.write("\t")
    TargetData.write(str(GrossReturnStd))
    TargetData.write("\t")
    TargetData.write(str(TradingQuintile))
    TargetData.write("\t")
    TargetData.write(str(FakeDummy))
```

```
        TargetData.write("\t")
        TargetData.write(SessionDummyEntry)

os.chdir(dir_input)
#Loop for finding relevant questionnaire-data for
    each main-entry
for questfilename in os.listdir(dir_input):
    #Open and split into data-points
    if (str(questfilename).startswith(str(filename
        [:11])) and str(questfilename).endswith('.
        txt')):
        QuestInputs = io.open(questfilename, 'r',
            encoding='utf-8').read().splitlines()
        ParticQuestEntry = QuestInputs[int(
            Participant)]
        QuestDataEntries = ParticQuestEntry.split('
            \t')

        #Creating two types of gender-dummies
        if QuestDataEntries[2] == "male":
            BinaryGenderDummy = "1"
            MaleDummy = 1
            FemaleDummy = "0"
            OtherGenderDummy = "0"
        else:
            if QuestDataEntries[2] == "female":
                BinaryGenderDummy = "0"
                MaleDummy = 0
                FemaleDummy = "1"
                OtherGenderDummy = "0"
            else:
                BinaryGenderDummy = "-"
                MaleDummy = 0
                FemaleDummy = "0"
                OtherGenderDummy = "1"
```

```
#Calculating BMI and two interpretative
    dummies
if not (QuestDataEntries[4]=="-" or
    QuestDataEntries[4]=="0" or
    QuestDataEntries[5]=="-" or
    QuestDataEntries[5]=="0"):
     BMI = 10000*int(QuestDataEntries[5])/
         float(math.pow(int(QuestDataEntries
         [4]), 2))
     if ( 18.5< BMI and 25 > BMI ):
          BMIDummyStrict = "0"
     else:
          BMIDummyStrict = "1"
     if ( 17< BMI and 30 > BMI ):
          BMIDummyExtreme = "0"
     else:
          BMIDummyExtreme = "1"
else:
     BMI = ""
     BMIDummyStrict = ""
     BMIDummyExtreme = ""

#Stock-Trading Experience Dummies
if QuestDataEntries[7] == "yes":
     StockExpDummyOne = 1
else:
     StockExpDummyOne = 0
if QuestDataEntries[8] == "yes":
     StockExpDummyTwo = 1
else:
     StockExpDummyTwo = 0

#Gambling Dummies
if QuestDataEntries[10] == "never":
```

```
        GamblingIndex = 0
    if QuestDataEntries[10] == "1-5 times":
        GamblingIndex = 1/float(3)
    if QuestDataEntries[10] == "6 - 10 times":
        GamblingIndex = 2/float(3)
    if QuestDataEntries[10] == "more than 10 
        times":
        GamblingIndex = 1

    #Dyscalculia-Dummies
    if QuestDataEntries[11] == "yes":
        DyscalculiaDummy = 1
    else:
        DyscalculiaDummy = 0

    #Creating Risk-Seeking Index and Dummy
        whether
    #statements are consistent with rationality
    RiskSeekIndex = 0
    previouselement = None
    ConsistencyHitCounter = 0
    RiskConsistDummy = "1"
    for element in QuestDataEntries[12:22]:
        RiskSeekIndex = RiskSeekIndex + int(
            element)
        if not previouselement == None:
            if not previouselement == element:
                ConsistencyHitCounter =
                    ConsistencyHitCounter + 1
        previouselement = element
    if ConsistencyHitCounter > 1:
        RiskConsistDummy = "0"
    RiskSeekIndex = RiskSeekIndex/float(10)

    #Creating Overconfidence-Index
```

```
ConfHitCounter = 0
ConfProblemDummy = 0
HighOverconfDummy = 0
if not (QuestDataEntries[22]=="-" or
    QuestDataEntries[23]=="-"):
     if not int(QuestDataEntries[22]) <= 56
        <= int(QuestDataEntries[23]):
          ConfHitCounter = ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[24]=="-" or
    QuestDataEntries[25]=="-"):
     if not int(QuestDataEntries[24]) <=
        6853 <= int(QuestDataEntries[25]):
          ConfHitCounter = ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[26]=="-" or
    QuestDataEntries[27]=="-"):
     if not int(QuestDataEntries[26]) <= 193
         <= int(QuestDataEntries[27]):
          ConfHitCounter = ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[28]=="-" or
    QuestDataEntries[29]=="-"):
     if not int(QuestDataEntries[28]) <= 24
        <= int(QuestDataEntries[29]):
          ConfHitCounter = ConfHitCounter + 1
else:
     ConfProblemDummy = 1
if not (QuestDataEntries[30]=="-" or
    QuestDataEntries[31]=="-"):
     if not int(QuestDataEntries[30]) <=
        380000 <= int(QuestDataEntries[31]):
```

```
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[32]=="-" or
    QuestDataEntries[33]=="-"):
    if not int(QuestDataEntries[32]) <=
        1756 <= int(QuestDataEntries[33]):
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[34]=="-" or
    QuestDataEntries[35]=="-"):
    if not int(QuestDataEntries[34]) <=
        6000 <= int(QuestDataEntries[36]):
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[36]=="-" or
    QuestDataEntries[37]=="-"):
    if not int(QuestDataEntries[36]) <=
        3474 <= int(QuestDataEntries[37]):
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[38]=="-" or
    QuestDataEntries[39]=="-"):
    if not int(QuestDataEntries[38]) <=
        9562 <= int(QuestDataEntries[39]):
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1
if not (QuestDataEntries[40]=="-" or
    QuestDataEntries[41]=="-"):
    if not int(QuestDataEntries[40]) <=
        10994 <= int(QuestDataEntries[41]):
```

```
        ConfHitCounter = ConfHitCounter + 1
else:
    ConfProblemDummy = 1

if ConfProblemDummy == 0:
    OverconfIndex = ConfHitCounter/float
        (10)
    if OverconfIndex > 0.7:
        HighOverconfDummy = 1
else:
    HighOverconfDummy = ""
    OverconfIndex = ""

#Creating Low-SelfesteemIndex
SelfEsteemSumOne = 15 - int(
    QuestDataEntries[42]) - int(
    QuestDataEntries[44]) - int(
    QuestDataEntries[45]) - int(
    QuestDataEntries[48]) - int(
    QuestDataEntries[51])
SelfEsteemSumTwo = int(QuestDataEntries
    [43]) + int(QuestDataEntries[46]) + int(
    QuestDataEntries[47]) + int(
    QuestDataEntries[49]) + int(
    QuestDataEntries[50])
LowSelfEsteemIndex = (SelfEsteemSumOne +
    SelfEsteemSumTwo)/float(30)


### Creating Personal/Treatment-
    Interactions ###

#Male Dummy
FeedTreatMaleDummyInter =
    FeedbackTreatDummy*MaleDummy
```

```
InforestTreatMaleDummyInter =
    InfoRestrTreatDummy*MaleDummy
TreatInterMaleDummyInter =
    TreatInteractDummy*MaleDummy

#Age
Age = float(QuestDataEntries[3])
FeedTreatAgeInter = FeedbackTreatDummy*Age
InforestTreatAgeInter = InfoRestrTreatDummy
    *Age
TreatInterAgeInter = TreatInteractDummy*Age

#Wealth
Wealth = float(QuestDataEntries[6])
FeedTreatWealthInter = FeedbackTreatDummy*
    Wealth
InforestTreatWealthInter =
    InfoRestrTreatDummy*Wealth
TreatInterWealthInter = TreatInteractDummy*
    Wealth

#Stockexperience One
FeedTreatStockExpOneInter =
    FeedbackTreatDummy*StockExpDummyOne
InforestTreatStockExpOneInter =
    InfoRestrTreatDummy*StockExpDummyOne
TreatInterStockExpOneInter =
    TreatInteractDummy*StockExpDummyOne

#Stockexperience Two
FeedTreatStockExpTwoInter =
    FeedbackTreatDummy*StockExpDummyTwo
InforestTreatStockExpTwoInter =
    InfoRestrTreatDummy*StockExpDummyTwo
TreatInterStockExpTwoInter =
```

```
    TreatInteractDummy*StockExpDummyTwo

#Speeding
Speeding = float(QuestDataEntries[9])
FeedTreatSpeedingInter = FeedbackTreatDummy
    *Speeding
InforestTreatSpeedingInter =
    InfoRestrTreatDummy*Speeding
TreatInterSpeedingInter =
    TreatInteractDummy*Speeding

#GamblingIndex
FeedTreatGamblIndexInter =
    FeedbackTreatDummy*GamblingIndex
InforestTreatGamblIndexInter =
    InfoRestrTreatDummy*GamblingIndex
TreatInterGamblIndexInter =
    TreatInteractDummy*GamblingIndex

#Dyscalculia
FeedTreatDyscalculiaInter =
    FeedbackTreatDummy*DyscalculiaDummy
InforestTreatDyscalculiaInter =
    InfoRestrTreatDummy*DyscalculiaDummy
TreatInterDyscalculiaInter =
    TreatInteractDummy*DyscalculiaDummy

#RiskSeekIndex
FeedTreatRiskSeekIndexInter =
    FeedbackTreatDummy*RiskSeekIndex
InforestTreatRiskSeekIndexInter =
    InfoRestrTreatDummy*RiskSeekIndex
TreatInterRiskSeekIndexInter =
    TreatInteractDummy*RiskSeekIndex
```

```
#Overconfidence Index
if not OverconfIndex == "":
    FeedTreatConfIndexInter =
        FeedbackTreatDummy*OverconfIndex
    InforestTreatConfIndexInter =
        InfoRestrTreatDummy*OverconfIndex
    TreatInterConfIndexInter =
        TreatInteractDummy*OverconfIndex
else:
    FeedTreatConfInter = ""
    InforestTreatConfInter = ""
    TreatInterConfInter = ""

#Overconfidence Dummy
if not HighOverconfDummy == "":
    FeedTreatConfDummyInter =
        FeedbackTreatDummy*HighOverconfDummy
    InforestTreatConfDummyInter =
        InfoRestrTreatDummy*
        HighOverconfDummy
    TreatInterConfDummyInter =
        TreatInteractDummy*HighOverconfDummy
else:
    FeedTreatConfDummyInter = ""
    InforestTreatConfDummyInter = ""
    TreatInterConfDummyInter = ""

#LowSelfEsteemIndex
FeedTreatLowSelfEstIndexInter =
    FeedbackTreatDummy*LowSelfEsteemIndex
InforestTreatLowSelfEstIndexInter =
    InfoRestrTreatDummy*LowSelfEsteemIndex
TreatInterLowSelfEstIndexInter =
    TreatInteractDummy*LowSelfEsteemIndex
```

```
#Creating Dummy to identify synthetic entries
FakeEntryDummy = 0
#if (str(OverconfIndex)=="1.0" and QuestDataEntries[32]=="0" and BMI=="-"):
if ((str(DataEntryList[0])=="171012_1723" and str(DataEntryList[4])=="10") or (str(DataEntryList[0])=="171115_1730" and str(DataEntryList[4])=="13") or (str(DataEntryList[0])=="171018_1718" and str(DataEntryList[4])=="28")):
    FakeEntryDummy = 1

#Writing everything down
os.chdir(dir_output)
with open(ResultName, 'a') as TargetData:
    TargetData.write(ParticQuestEntry)

with open(ResultName, 'a') as TargetData:
    TargetData.write("\t")
    TargetData.write(BinaryGenderDummy)
    TargetData.write("\t")
    TargetData.write(str(MaleDummy))
    TargetData.write("\t")
    TargetData.write(FemaleDummy)
    TargetData.write("\t")
    TargetData.write(OtherGenderDummy)
    TargetData.write("\t")
    TargetData.write(str(BMI))
    TargetData.write("\t")
    TargetData.write(BMIDummyStrict)
    TargetData.write("\t")
    TargetData.write(BMIDummyExtreme)
    TargetData.write("\t")
    TargetData.write(str(StockExpDummyOne))
```

```
TargetData.write("\t")
TargetData.write(str(StockExpDummyTwo))
TargetData.write("\t")
TargetData.write(str(GamblingIndex))
TargetData.write("\t")
TargetData.write(str(DyscalculiaDummy))
TargetData.write("\t")
TargetData.write(str(RiskSeekIndex))
TargetData.write("\t")
TargetData.write(RiskConsistDummy)
TargetData.write("\t")
TargetData.write(str(OverconfIndex))
TargetData.write("\t")
TargetData.write(str(HighOverconfDummy)
    )
TargetData.write("\t")
TargetData.write(str(LowSelfEsteemIndex
    ))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterMaleDummyInter))
TargetData.write("\t")
TargetData.write(str(FeedTreatAgeInter)
    )
TargetData.write("\t")
TargetData.write(str(
    InforestTreatAgeInter))
TargetData.write("\t")
TargetData.write(str(TreatInterAgeInter
```

```
    ))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatWealthInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatWealthInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterWealthInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterStockExpOneInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterStockExpTwoInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatSpeedingInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatSpeedingInter))
TargetData.write("\t")
```

```
TargetData.write(str(
    TreatInterSpeedingInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatGamblIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatGamblIndexInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterGamblIndexInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterDyscalculiaInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    TreatInterRiskSeekIndexInter))
TargetData.write("\t")
TargetData.write(str(
    FeedTreatConfIndexInter))
TargetData.write("\t")
TargetData.write(str(
    InforestTreatConfIndexInter))
```

```
        TargetData.write("\t")
        TargetData.write(str(
            TreatInterConfIndexInter))
        TargetData.write("\t")
        TargetData.write(str(
            FeedTreatConfDummyInter))
        TargetData.write("\t")
        TargetData.write(str(
            InforestTreatConfDummyInter))
        TargetData.write("\t")
        TargetData.write(str(
            TreatInterConfDummyInter))
        TargetData.write("\t")
        TargetData.write(str(
            FeedTreatLowSelfEstIndexInter))
        TargetData.write("\t")
        TargetData.write(str(
            InforestTreatLowSelfEstIndexInter))
        TargetData.write("\t")
        TargetData.write(str(
            TreatInterLowSelfEstIndexInter))
        TargetData.write("\t")
        TargetData.write(str(FakeEntryDummy))

#Next Line in Data
with open(ResultName, 'a') as TargetData:
    TargetData.write("\n")
```

Chapter 3

Growth Diagnostics: A Foundation

co-authored with Sayantan Ghosal

3.1 Introduction

Taking as its starting point the divergence between private and social valuations of economic activity, growth diagnostics (Hausmann et al., 2005a, 2008; Rodrik, 2010) is a strategy for identifying the priorities for policy reform in settings characterized by second-best distortions (Pigou, 1954; Lipsey and Lancaster, 1956). As Hausmann et al. (2005a) put it: "The idea behind the strategy is simple: if (a) for whatever reason the full list of requisite reforms is unknowable or impractical, and (b) figuring out the second-best interactions across markets is a near-impossible task, the best approach is to focus on the reforms where the direct effects can be reasonably guessed to be large. ... The principle to follow is simple: go for the reforms that alleviate the most binding constraints, and hence produce the biggest bang for the reform buck." (p. 7)

A vast policy literature uses growth diagnostics as a policy framework, from the early work by the World Bank (World Bank, 2005, with a focus on emerging market economies and eastern European Economies e.g. Cambodia, India, Brazil, the Baltic States) and the Asian Development Bank (Felipe et al., 2011), to the more recent initiative by the Scottish Government (Scottish Government, 2016) focusing on inclusive growth and the related work in New Zealand (Karacaoglu, 2015). The bulk of this literature uses growth diagnostics as a tool for identifying priorities for policy reform at a given point in time.

However, the empirical literature on policy reform and growth (Hausmann et al., 2005b) suggests that policy reform is an adaptive process that occurs over a period of time; moreover, policy reforms that work in some settings do not necessarily work in all settings (Rodrik, 2008), conventional wisdom is often wrong (Campos and Coricelli, 2002), the sequencing of policy reform matters and policy-makers may get the sequencing wrong (Edwards, 1990).

Further, second-best distortions may interact in ways that mean the indirect social welfare losses of relaxing a second-best constraint swamp the direct social welfare gain from doing so (see Bergoeing et al., 2015, for an empirical analysis in the context of firm entry/exit and adoption of new technologies in emerging market economies). Reducing a distortion in one activity introduces distortions in activities that, before the policy reform, were otherwise undistorted e.g. mitigating the underprovision of a public good by a distortionary income tax. If the policy-maker wrongly estimates the benefits and costs involved, there may be an overall net social welfare loss.

Motivated by the preceding considerations, this paper provides a formal analysis of the conditions under which an adaptive implementation of growth diagnostics converges to a socially desirable growth outcome even allowing for the possibility of mistakes i.e. situations where the losses dominate the gains.

While the policy-maker may have a reasonably accurate estimate of the first-order approximation (direct effect) of reducing a distortion, information about the second-order effects of doing so is less reliable. Nevertheless, the policy-maker may have access to different sources of information (signals generated by an unknown stochastic process) that produce estimates of potential losses to social welfare from the second-order interaction effects. When these losses outweigh the first-order welfare gains, a mistake occurs.

We study an adaptive policy process, which, under minimal assumptions on the nature of uncertainty ensures (Proposition 1) that realized gross social welfare losses are almost surely (with probability one) bounded. This condition, in turn, implies convergence to the socially optimal outcome. The adaptive policy process ensures that the change in the weight assigned to different signals point in the same direction as net social welfare gains irrespective of the outcome, a property similar to one used in the proof of Blackwell's approachability theorem (Blackwell, 1956). Then, Proposition 2 shows that *any* policy process (and not just the one formally analysed below in section 3.2) which satisfies the condition that losses are almost surely bounded will converge to the socially optimal outcome.

How economic insights should be translated into governmental legislation is a long-standing debate. Lindblom (1959, 1979) was an early proponent of incremental experimentation in policy-making. He argued that any economic analysis comes with a high degree of uncertainty and must be regarded as incomplete. In a static policy-setting, he showed that reforms based on these analyses should reflect the underlying incompleteness by implementing changes in small steps and profiting from the information feedback rather than reaching for swiping "all-or-nothing" solutions. While Lindblom (1959, 1979) saw *muddling* as the best of a number of bad options, scholars have since gone beyond these claims in examining the role of explicit and purposeful experimentation in policy-making (Hayek, 1978; North, 1990; Roland, 2000; Mukand and Rodrik, 2005). Our theoretical results contribute to this literature by showing how a process of policy making involving adaptive experimentation can, under minimal assumptions on the uncertainty involved, converge to the socially optimal outcome. In Section 4, we discuss a number of different historical examples to illustrate how such a process could play out in practice.

Our results are related to the convergence properties of iterative processes in public good and collective consumption by Dreze and de la Vallee Poussin (1971), Malinvaud (1972) and retrading in market games (Ghosal and Morelli, 2004) where reallocations can be Pareto improving at each step. However, in contrast to Proposition 2, these papers do not allow for the possibility of mistakes. Ghosal and Porter (2013) in their analysis of cautious retrading in the context of exchange economies assume a property related to the condition that losses are almost surely bounded; however, in contrast to Proposition 1, they do not derive it as an outcome of an explicitly defined adaptive forecasting and policy implementation process. The formulation the adaptive policy process and the first part of the proof of Proposition 1 is related to Cesa-Bianchi and Lugosi (2004).

In specific settings, taking into account both the direct and indirect effects of policy reform that target second-best distortions is particularly important with limited fiscal resources, a point emphasized by Martin and Pindyck (2015) and in policy discussions by the New Zealand Government (Karacaoglu, 2015, 2016) and various Scottish Government papers linked to the use of growth diagnostics (e.g. Gillespie, 2016; Scottish Government, 2016).

The remainder of the paper is structured as follows. Section 2 sets out a reduced form second-best policy framework; section 3 is devoted to stating and deriving the main results of the paper. Section 4 is devoted to a discussion of relevant historical examples. The last section concludes.

3.2 A Reduced-Form Policy Framework

We begin by setting out a reduced form policy framework that takes as the starting point the role of second-best distortions which introduce a wedge between social and private marginal valuations of economic activity. For the purpose of illustration, we accompany the general introduction of the model with a simplified, stylised example.

Let $X = \{X_{i,a,t} : i \in I, a \in A, t \in T\}$ denote a trajectory where I denotes a list of activities, A denotes a list of economic agents and T (where T could be finite or infinite) denotes a list of time periods. For our stylised example, we will consider an economy with two economic agents, so that $A = \{1, 2\}$. There are two activities those agents may follow, meaning $I = \{1, 2\}$. On the one hand, agents may drive environmentally friendly cars. Alternatively, agents may, for leisure, drive large and fast cars that are harmful to the environment. Time, T, is infinite.

Let $\mathbf{X} \subset \Re^{IAT}$ denote the set of feasible trajectories. Feasibility, in this context, would include constraints such as intertemporal budget (resource) and technological constraints. We allow for the feasible state to depend on the initial state summarizing the history of the economy. We assume that $\mathbf{X}$ is a compact, convex set with a non-empty interior. [1]

Let $W : \mathbf{X} \rightarrow R$ denote a continuous, strictly concave twice continuously (Frechet) differentiable social welfare function that has a unique maximum (denoted by Y) in the interior of $\mathbf{X}$.

Let $X \neq Y$ denote the equilibrium trajectory generated by some underlying pattern of market/strategic interactions between agents over time. As $X \neq Y$, it must be the case that there are distortions which introduce a wedge between social and private marginal valuations in economic activity. Let $u_a : \mathbf{X} \rightarrow R$ denote a continuous, concave, twice continuously (Frechet) differentiable utility function that represents the preferences of agent a over the set of feasible trajectories.

Let $u_{a,i,t}(X)$ be the derivative of u_a to the i^{th} and t^{th} element and, analogously, $W_{i,t}$ the derivative of $W(X)$ to the i^{th} and t^{th} element. Thus, $u_{a,i,t}(X)$ denotes the marginal valuation of activity i by agent a at period t.[2] It follows that for at least one activity i and agent a, there exists a constant $d_{a,i,t} \neq 0$ such that

$$W_{i,t}(X) - u_{a,i,t}(X) - d_{a,i,t} = 0, i \in I, a \in A, t \in T \qquad (1)$$

[1] Compactness is preserved from $X_{i,a,t}$.

[2] As both $u_{a,i,t}(X)$ and $W_{i,t}$ are marginal concepts, separability of preferences is not required.

where along trajectory X, $W_{i,t}(X)$ denotes the social valuation of activity i at time t, $u_{a,i,t}(X)$ denotes the corresponding private valuation by agent a and $d_{a,i,t} \neq 0$ for some $i \in I, a \in A, t \in T$ is the associated distortion.

For the stylised example, we assume that the environmental impact of driving environmentally friendly cars is small. As an agent carrying out this activity imposes weak negative effects on other agents, we state that an agent's private valuation of this activity differs slightly from the social valuation of it. On the other hand, driving environmentally harmful cars causes pollution and thereby decreases the welfare of other agents strongly. Therefore, there exists a strong for potential of the social valuation of driving environmentally harmful cars to differ from the private valuation. we further assume that agent 1 has an inherent disposition to act in an environmentally conscious fashion so that the agent never engages is in driving harmful cars, but instead solely drives environmentally friendly ones. Agent 2, however, does not have this disposition and instead engages in driving environmentally harmful cars.

Consider, a second-best social welfare maximization where at a given point in time t, taking the trajectory X as a starting point, the social planner maximizes W over $\mathbf{X}$ treating (1) as a system of constraints. Let $d_t = (d_{a,i,t} : i \in I, a \in A)$. Let $\widetilde{W}(d_t)$ denote the value function of the maximization problem viewed as a function of d_t. Consider a vector of perturbations in a small neighbourhood of zero, $\varepsilon_t = (\varepsilon_{a,i,t} : i \in I, a \in A)$. Suppose the policy-maker attempts to alter the second-best distortions d_t by ε_t. Using a (second-order) Taylor series expansion, we obtain:

$$\begin{aligned}\Delta_{\varepsilon_t}(X; d_t) &= \widetilde{W}(d_t + \varepsilon_t) - \widetilde{W}(d_t) \\ &= \sum_{a,i} \varepsilon_{a,i,t} \frac{\partial \widetilde{W}(d_t)}{\partial d_{a,i,t}} + \frac{1}{2} \sum_{a,i} \sum_{a',i'} \varepsilon_{a,i,t} \varepsilon_{a',i',t} \frac{\partial^2 \widetilde{W}(d_t)}{\partial d_{a',i',t} \partial d_{a,i,t}} \quad (2)\end{aligned}$$

Varying ε_t can be seen as implementing policy-packages at time t. For the accompanying example, we will consider two single policies aimed at cutting back the socially detrimental overindulgence in driving cars. First, we consider a tax on gasoline. Second we consider a toll levied for the usage of public roads.

For each $a \in A$, and $i \in I$ and $t \in T$, let $\mu_{a,i,t}$ denote the Lagrangian multiplier associated with (1) in the maximization problem where taking the trajectory X as a starting point, the social planner maximizes W over $\mathbf{X}$ treating (1) as a system of constraints. Clearly $\mu_{a,i,t}$ can be interpreted as the marginal (incremental) social value along the trajectory X of reducing the distortion of activity i for agent a at time t. As

$\widetilde{W}_t(d_t)$ is a value function, using duality, it follows that:

$$\frac{\partial \widetilde{W}_t\left(d_t\right)}{\partial d_{a,i,t}} = -\mu_{a,i,t} \text{ for all } a, i, t \text{ such that } d_{a,i,t} \neq 0 \text{ (3).}$$

The first term in equation (2) captures the *direct effect* on social welfare, at a given point of time, when distortion of activity i' for agent a' at time t' is reduced. Using equation (3), it is clear that when activity i is undistorted so that $d_{a,i,t} = 0$, the social valuation and private valuation of activity i at time t are identical. Along a specific trajectory X, the more socially costly is the distortion in activity i for agent a at time t, the higher is the value of $\mu_{a,i,t}$, the current marginal social valuation of reducing the distortion $d_{a,i,t}$.

The second term in equation (2) captures the *indirect effect* on social welfare of reducing the distortion $d_{a,i,t}$ via its impact on the marginal social valuation of reducing the distortion valuation $d_{a',i',t}$: when there are other distorted activities, on the margin, the interaction effects between the different distortions will matter. The sum of the second-order indirect effects captures the effect of changing the distortion in activity i' for agent a' at time t' on the marginal social value of reducing distortions in all other activities i for other agents a.

For the considered policies of taxing gasoline and collecting tolls for the usage of public roads, it may easily be imagined that both policies will individually have a negative impact on the valuation for driving all kinds of cars for all agents. Moreover, as all agents privately overvalue all manners of car-driving, we assume that the package of implementing both policies, via extra costs for all agents carrying out the considered activities, would work towards realigning private valuations with social valuations in terms of the *direct effects* of both policies. The same holds true for both policies individually.

However, both policies, if implemented simultaneously, may also interact with each other in changing the agents' private valuations of the considered activities. In particular, although implementing tolls for public roads when a tax on gasoline is already in place leads to an added positive *direct* effect in aligning private valuations of driving environmentally-friendly cars, both effects could reinforce each other to the extent that the private valuation of driving environmentally friendly cars dips below the social valuation. While before implementing both policies one would expect overindulgence in driving environmentally friendly cars, one would expect underindulgence after it, due to undesirable *indirect effects* from the policy package. In fact, we assume that such negative indirect effects outweigh the direct beneficial effects of both individual policies, leading to an overall decrease in social welfare when the policy-package is implemented. In this case, it would be in the social planner's interest to not implement tolls for public roads in the presence of a tex of gasoline despite their beneficial direct effects on social welfare.

3.3 Adaptation, Mistakes and Convergence to Social Optimum

Heuristically, growth diagnostics can be viewed as an adaptive strategy for identifying the priorities for policy reform. In this section, we model growth diagnostics as an adaptive policy tool i.e. at each point in time, given the prevailing trajectory of the economy, reduce the second-best distortion or gap with the greatest direct social impact. However, in attempting to do so, the policy maker may not have enough information to rule out losses in social welfare from an application of the growth diagnostic. What are the conditions under which an adaptive application of growth diagnostics results in a sequence of changes in social welfare that converge to the socially optimal outcome even allowing for the possibility of mistakes? In this section, we model an adaptive forecasting and policy implementation process, which under minimal assumptions on the nature of uncertainty over potential welfare losses, converges to the socially optimal allocation.

Starting from some feasible trajectory $X \neq Y$, at a given t and the associated vector of distortions d_t, the policy-maker implements $d_t+\varepsilon_t$, where $\varepsilon_t = \arg\max_{\varepsilon_t \in N(0)} \Delta_{\varepsilon_t}(X; d_t)$ (where $N(0)$ is a small neighbourhood of zero) with the objective of achieving gross social welfare gains $\Delta(X; d_t)$ (upto a second-order approximation). For later reference, we label this as **implementing growth diagnostics** at time t.

For the stylised example, we will restrict our attention to a limited number of actions, or policy-packages, which the social planner can introduce. The social planner can decide to only tax gasoline, only levy tolls for public roads or implement both measures simultaneously. While all of these policy-packages can be introduced, they can also be revoked in any combination. We assume that the ordering of all considered policy-packages in terms of social welfare is as follows, best to worst: taxing gasoline, levying tolls on public roads, no policies, taxing gasoline and levying tolls.

Starting at some initial time period $t = 0$, an **adaptive implementation of growth diagnostics** generates a sequence of trajectories $\left(\widehat{X}(t) : t \geq 0\right)$ and an associated sequence of (potential) gross social welfare gains $\left(\Delta\left(\widehat{X}(t); d_t\right) : t \geq 1\right)$.

Starting from some trajectory X at time t, let $b_t \geq 0$ denote the positive effects (benefits) on social welfare incurred after implementing $d_t + \varepsilon_t$ and let $l_t \geq 0$ denote the negative effects (loss) on social welfare welfare incurred after implementing $d_t + \varepsilon_t$. The net social welfare gains at time t are then $b_t - l_t$ (upto a second-order approximation). We say that the policy maker makes a **mistake** while implementing the growth diagnostics whenever $b_t - l_t < 0$ i.e. there is a net social welfare loss at time t. Framed in the context of the stylised example, introducing both a tax on gasoline as wells as tolls for public roads would constitute such a mistake. Negative indirect interaction effects (see section 3.2 for a further discussion) of the policy-packages swamp any direct benefits of both individual policies and thus lead to an overall decrease of social welfare.

We assume that $l_t \leq \tilde{L}$, $\widetilde{L} > 0$ at each t where the bound $\tilde{L}$ on the gross social welfare loss in any one period is uniform over $\mathbf{X}$. Furthermore, assume that there are $s = 1, ..., N$ sources of signals about the gross social welfare loss. Let $l_{s,t} \leq \widetilde{L}$ denote expected gross social welfare losses given signal s that could be incurred when growth diagnostics is implemented at time t. For each signal $s = 1, ..., N$, the policy-maker computes the net social welfare $b_0 - l_{s,0}$.

Assume that the policy maker puts a weight $f_{s,0} = 1$ on each signal s. Let $f_{s,t}$ denote the weight on signal s at some $t \geq 0$. Given α, $\frac{1}{2} \leq \alpha < 1$, if there are at least as many signals as αN such that $b_0 - l_{s,0} \geq 0$ (a net social welfare gain), growth diagnostics is implemented at $t = 0$; otherwise, the policy maker does nothing and the economies continues along $X(0)$ for one more time period.

If a mistake occurs (i.e. there is an actual net social welfare loss) the policy-maker reduces the weight on all signals s such that $b_0 - l_{s,0} \geq 0$ by a fixed fraction β, $0 < \beta < 1$.

In any subsequent period $t > 0$, the policy maker compares the total weight of the experts predicting a net social welfare gain with those predicting a net social welfare loss at $X(t)$ and growth diagnostics is implemented only if the total weight on signals predicting a net social welfare gain exceeds αF_t, where F_t is the total weight on signals in period t; otherwise, the policy maker does nothing and the economy continues along $X(t)$ for one more time period.

As before, at each t where a mistake occurs the policy-maker reduces the weight on all signals s such that $b_t - l_{s,t} \geq 0$ by a fixed fraction β, $0 < \beta < 1$.

We make one assumption about the underlying stochastic processes generating the signals, **the almost surely finite mistakes property**: there exists a signal s^* (the policy-maker does not know which one) such that, with probability one, along the sequence of trajectories $\{X(t) : t \geq 0\}$ generated by an adaptive implementation of growth diagnostics, the corresponding sequence $\{l_t : t \geq 0\}$ is such that $b_t - l_{s^*,t} < 0$ only for finitely many t, $t \geq 0$: this number (assumed to be the minimum over all signals which satisfy the above property) is denoted by m^{*}.

We make no other assumptions about the underlying stochastic processes generating the signals, specifically, we do not assume that the policy maker knows that the underlying stochastic processes generating the signals has the almost surely finite mistakes property.

Note that, by construction, the adaptive forecasting and policy making process satisfies the following condition for each $t \geq 0$:

$$\sum_{s} \left[\Delta\left(X(t)\right) - l_t\right]\left[f_{s,t+1} - f_{s,t}\right] \geq 0 \ (4)$$

where the weight $f_{s,t}$ for each source s changes according to whether a false prediction has been made or not. In words, the change in the weight assigned to different signals point in the same direction as net social welfare gains, irrespective of the outcome. This condition is similar to a key property used in the proof of Blackwell's approachability theorem (Blackwell, 1956).

Let $l_t \in \{l_{s,t} : s = 1, ..., N\}$ denote the realised gross social welfare loss incurred at time t when growth diagnostics is implemented at time t. We say that **losses are almost surely bounded** if $\sum_{t=0}^{\infty} l_t \leq L$, for some finite L, with probability one.

In order to illustrate how the given learning environment affects the process of policy-making, we will now lay out a number of representative cases making use of the stylised example we have introduced so far. Note that in all cases the socially optimal solution is reached if gasoline is taxed but no tolls for public roads are levied. In all cases, we will direct special attention to potential obstacles in achieving this outcome and how the adaptive policy process described can overcome them.

For simplicity, we choose two sources of signals, called 'experts', providing information to the social planner on projected welfare losses associated with any given policy-package while $\alpha = 0.5$. Hence, whenever the social planner assigns the same weight to both experts, a single recommendation from any of the experts is sufficient for the social planner to implement a policy-package. On the other hand, discrediting any of the two experts by deflating the associated weight by any $0 < \beta < 1$ as the result of a false recommendation will lead to a clear favour towards the other expert.

It should be noted, that the model always allows for the possibility of inaction in any given given time period so that no policies are implemented or revoked. This occurs if none of the policy-packages the social planner contemplates gathers a critical level of support, i.e. at least one expert projecting a net welfare gain as a result of the change. This case can be interpreted as experts debating among themselves as well as potential uncertainty about the future being resolved as time passes. As the more interesting cases, from a modelling perspective, are time periods where actions are taken, we will focus on these.

Case 1: no possibility for mistakes

We will first consider a case where the projected welfare losses by the two available experts are always correct on all policy-packages the social planner considers. For the chosen setting, the planner needs at least one expert to support (via the stated loss estimate) any given policy-package in order to implement it. However, if all experts always provide correct information on all packages, a policy-package leading to a welfare loss can never garner any support. Therefore, the social planner cannot make mistakes in this scenario. Moreover, all policy-packages that do, in fact, entail net social welfare gains will always gather full support and thus will be implemented. For an economy starting on a sub-optimal trajectory in the given setting, this represents a best-case scenario. The policy-maker will implement the correct policy-mixture (taxing gasoline but not levying tolls on roads) while not incurring any welfare losses from policy-mistakes.

Case 2: temporary possibility for mistakes

Next, we will consider the case where one of the experts may provide incorrect projections on welfare losses of policy-packages to the social planner. In particular, we assume that one of the experts projects that implementing both a tax on gasoline as well as levying tolls on public will not lead to any indirect interaction effects between the two single policies. As the direct effects of both individual policies are set to be positive, the social planner will be tempted to implement the policy-package. Since also at least one expert predicts net social welfare gains as a result of this policy, the social planner may indeed implement it. In the terminology of our model, this constitutes a mistake while implementing growth diagnostics.

While one of the experts provides incorrect advice, we assume in this scenario that the other expert always provides correct advice. The expert providing correct advice thus satisfies the almost surely finite mistakes property with the addition that, in this particular case, there is no intermediate time span where the expert in question may provide incorrect information.

After committing a mistake, the social planner will review actual changes in social welfare and adjust the weight attached to all experts according to their previous projections. The weight attached to the expert providing incorrect projections will thus be deflated by β. The weight of the other expert, on the other hand, will remain unchanged, as this expert will have correctly predicted a net welfare loss. As a result of this adaptation, the expert always giving correct advice will from there on always surpass the critical threshold of $\alpha = 0.5$, whereas the expert giving incorrect advice never will.

At this point, the economy will have effectively transitioned into Case 1 described above where the social planner will never receive any false information that can affect the actions. Welfare losses from policy-making mistakes are temporary and, as welfare losses in any given period are assumed to be bounded, they are bounded over the duration of the process of implementing growth diagnostics.

Case 3: prolonged possibility for mistakes

Last, we consider the case, where, starting from a given time period, both experts may provide incorrect projections on the welfare losses associated with policy changes. In particular, we now assume that both experts consistently predict no indirect interaction effects on social welfare between the two policies of taxing gasoline and levying tolls on public roads.

As in the previous case, due to the positive direct effects of the policy- package as well as the critical support from experts, the social planner will implement the policy-package and, thus, make a mistake. However, as now both experts incorrectly predicted net welfare gains, the social planner will deflate both of their weights by β. As long as the experts do not alter their prediction patterns, the economy may enter a loop of revoking and implementing the policy-package in alternating periods. Even though the overall weight on experts will keep declining, the detrimental policy- package will not lose critical support. In this scenario, the social planner may therefore be prone to enter a prolonged period of mistakes, leading to social welfare losses accumulating over time.

However, the almost surely finite mistakes property ensures that such losses may not be continue indefinitely. By construction, one expert must begin solely providing correct projections of welfare losses associated with policy changes after a finite amount of time. This has two main effects. First, although mistakes may still be conducted initially, experts providing support for them will lose critical support (in relative terms). Hence, through the social planner adapting weights, the propensity for committing further mistakes will be curbed. Second, policy changes actually leading to net social welfare gains are now ensured to be proposed. Again, the mechanism behind the weight adjustments of experts will ultimately favour those policies.

Qualitatively, the economy will have transitioned into Case 2 described above. Despite the possibility for continuous mistakes compounding losses in terms of social welfare, these are bounded over the duration of the optimisation problem as by the almost surely finite mistakes property one of the experts will start providing solely correct estimates after a *finite* amount of time and losses in any given period are assumed to be bounded.

These simplistic examples illustrate that under a variety of key scenarios, implementing growth diagnostics will lead to social welfare losses being bounded. These intuitions, indeed, maintain in a general setting (see Appendix E for a proof).

Proposition 1: *Consider the adaptive policy process described above. Then, losses are almost surely bounded.*

This insight proves crucial for determining the economy's welfare. The next proposition demonstrates that *any* process which satisfies the condition that losses are almost surely bounded will converge to the socially optimal outcome (Proof in Appendix E):

Proposition 2: *When losses are almost surely bounded, starting from a socially suboptimal trajectory X, the adaptive forecasting and policy implementation process leads to the socially optimal trajectory Y.*

3.4 Discussion

The fundamental insight form the theoretical analysis is that an adaptive forecasting and policy-making process where small adjustments are made in the direction of potential welfare gains will over time converge to a socially optimal outcome even allowing for the possibility of making mistakes. In this section, we illustrate how such a process could play out in practice with reference to a number of actual historical examples.

The potential benefits of incremental, adaptive policy experimentation, as well as the dangers of not doing so, are well exemplified by the economic history of China. In 1958, the communist government launched the Great Leap Forward movement. Its goal was to boost the country's industrial production. Ambitions for the project went as high as surpassing the United Kingdom in terms of raw output within 15 years. In order to achieve these dramatic results, resources, in the form of labour reallocation and increased direct taxation, were diverted from the Chinese agriculture (among other sectors).

In the short-run, the reform successfully achieved some its intended direct effects by raising the steel output from 5.35 million tons in 1957 to 18.66 million tons in 1960. However, the estimated capabilities of the agricultural sector to cope with increased quotas with fewer dedicated inputs turned out horrendously false.

The outcome was a disaster. (Li and Yang, 2005) provide a detailed analysis of the events. China's food production collapsed, with grain output falling by about 15% in each of the years 1959-1961. Even by the most conservative estimates, the resulting famine claimed the highest number of lives in recorded world history. The country's gross domestic product stagnated in 1960 and then fell by 16% and 5.6% in 1961 and 1962, respectively the two worst performances from 1952 until today (data from National Bureau of Statistics of China, 2017). Li and Yang (2005) show that the communist party's proclaimed explanation of bad weather as the main cause of the events accounts for only 13% of food shortages. They found that over 60% of the decline in grain production, and thus the ensuing humanitarian and economic crisis, was attributable to errors in decision-making of the government.[3]

The development strategy of the economically highly successful post-Mao era contrasted starkly from its precursor. Deng Xiaoping famously coined the Chinese style of policy-making of the time as "crossing the river by feeling the stones"- a gradual, adaptive and directed progression that preceded each step in reforming by careful testing. This approach was in part necessitated by the lack of the ruling elite of an exact and comprehensive blueprint for China's future (Lin et al., 2003; Naughton, 2006). But rather than forcing the intended changes through the political opposition, the adaptive reform implementation took the existing uncertainties deliberately into account. In fact, Heilmann (2008a,b) gives policy experimentation a front-row seat in China's economic miracle.

Many times over, momentous and potentially risky policy changes in China's transformation process were first tried out and studied on isolated parts of the economy before they were scaled-up. Long before China joined the World Trade Organisation and committed to today's level of openness to foreign enterprises, the government piloted special economic zones: geographically clearly defined areas in which the effects of internationalisation on the local economies could be observed and evaluated. The transition from a planning economy with centrally fixed prices to a market economy was mediated via a dual-price system: quotas had to be delivered at plan-prices while any surplus could be sold for private profits by the producers.

[3]Another example is the desiccation of the Aral Sea under the rule of the former Soviet Union. Starting in the 1940s, the water inflows of the then fourth largest lake in the world were diverted for irrigation in order to foster agricultural production in the surrounding desert environment. While the project achieved the desired direct effects by temporarily making Uzbekistan the largest exporter of cotton in 1980 (National Cotton Council of America, 2017), the environmental and economic repercussions were vastly underestimated. By 1988, the continued shrinking of the lake and its increasing salination already caused losses of billions of rubles by entirely wiping out the area's fishing industry, damaging other branches and causing serious health problems in the population (Micklin, 1988). By 2007, the lake's surface area had shrunk to 10% of its original surface area with recent counter-measures only slowly reversing the trend (Micklin and Aladin, 2008).

More recently, China continues to pursue the internationalisation of its currency, the Renminbi, in a similar fashion. In 2004, the government showed first signs of its intentions by allowing Hong Kong based banks to open Renminbi accounts with limited conversion allowances. In 2009, the Renminbi was legitimised to settle trade with ASEAN members as well as Macao and Hong Kong in five Chinese cities. Upon satisfactory results, the scheme was first extended in mid-2010 so that firms in twenty provinces could settle all trade in the local currency. Nowadays essentially all trade settlements may be executed with it (Eichengreen and Kawai, 2015).

While the economic history of China tells a compelling story of the dangers of "Big-Bang reformism" and the potential benefits of deliberate experimentation, both instances are by no means purely Chinese phenomena. Mukand and Rodrik (2005) note that while most Latin American countries vigorously underwent structural reforms, the national governments forced the commonly agreed best practices at the time upon the national economies with little care given to local needs. As a result, nearly all those countries achieved slower growth rates after the reforms than before.

A growing number of governments around the world reflect the lessons learned from history by harnessing the power of adaptive, policy-experimentation. The United Kingdom is at the forefront of the movement. Jowell (2003) summarises the already extensive role of pilot-studies in the UK's policy-making in the early 2000s. The national trend lately cumulated into the formation of the What Works Network with a total of nine institutions devoted to supplement political decision-makers with field evidence at crucial stages of reform design.

Finally, List and Gneezy (2014) cite the example of the movie delivery platform Netflix to argue for policy-experimentation not only on the side of national governments, but also as the cutting-edge in business managements. In 2011, the management of the then rising star in the market made a series of unfortunate decisions which resulted in falling customer subscriptions and subsequently plummeting stocks that brought the firm to the brink of bankruptcy. Today, literally every product change at the company is first tried out on an isolated cell of the consumer base. As the reason for their strategy the company state that such interventions are simply "too risky to roll out without extensive (...) testing." (Techblog, 2017) While these trials, in contrast to the policy-changes discussed in our main model, are carried out only on subsets of the relevant populations and are thus comparatively sterile in character, they exemplify the importance of using results of policy impact as feedback to guide reform priorities in the future.

3.5 Conclusions

In this paper, we view growth diagnostics as an iterative tool for identifying the priorities of policy reform: at each point in time, reduce the section-best distortion or gap with the greatest direct social impact. We show that such a process converges to the socially optimal allocation even if the policy-maker incurs unintended social welfare losses.

The formal analysis reported here provides a set of convergence results that demonstrate that growth diagnostics is a robust, flexible policy framework. However, in practice, policy reform is constrained by political interests. The results we obtain here constitute a starting point for an analysis that take such political constraints explicitly into account.

In future work, we plan to extend the analysis reported here to examine how special interest groups can attempt to influence which second-best distortions are addressed and the impact such influence can have on probability of making mistakes.

Appendix E

Proofs

E.1 Proof of Proposition 1

Starting from some trajectory $X \neq Y$, at each $t \geq 0$, there are two possibilities: (1) the total weight on signals predicting a net social welfare gain exceeds αF_t, and (2) the total weight on signals predicting a net social welfare gain is strictly less than αF_t. At t, a mistake is only possible if (1) prevails.

Let F_m denote the total weight on signals when the policy-maker makes the m-th mistake. At the time period when the policy maker makes the m-th mistake, the overall weight of the signals forecasting a net social welfare gain must be at least αF_m; the weight on the signals forecasting a net social welfare loss is unchanged at $(1-\alpha)\,\beta F_{m-1}$.

By computation it follows that

$$\begin{aligned} F_m &\leq (1-\alpha)\,\beta F_{m-1} + \alpha F_{m-1} \\ &= F_{m-1}\left((1-\alpha)\,\beta + \alpha\right). \end{aligned}$$

Moreover, $F_m \geq \beta^{m^*}$ as the current weight on signal s^*, $f_{s^*,t}$, must be such that $f_{s^*,t} \geq \beta^{m^*}$ (given that $0 < \beta < 1$) and $F_m \geq f_{s^*,t}$. Therefore, it must be the case that

$$\beta^{m^*} \leq F_{m-1}\left((1-\alpha)\,\beta + \alpha\right)$$

from which it follows that

$$\beta^{m^*} \leq F_0\left((1-\alpha)\,\beta + \alpha\right)^m$$

which, by computation, is equivalent to

$$\left(\frac{1}{(1-\alpha)\,\beta+\alpha}\right)^{m} \leq N\left(\frac{1}{\beta}\right)^{m^*}.$$

Furthermore,

$$m \ln\left(\frac{1}{(1-\alpha)\,\beta+\alpha}\right) \leq \ln N + m^* \ln\left(\frac{1}{\beta}\right)$$

which, by computation, is equivalent to

$$m \leq m_{\max} = \frac{\ln N + m^* \ln\left(\frac{1}{\beta}\right)}{\ln\left(\frac{1}{(1-\alpha)\beta+\alpha}\right)}$$

as $(1-\alpha)\,\beta+\alpha<1$.

Therefore, along the sequence generated by the adaptive policy process, with probability one, there is at most a finite number of mistakes and the sum over gross realised social welfare losses across all the time periods in which mistakes are made is bounded above by $m_{\max}\widetilde{L}$.

When a mistake *does not* occur at time t, we have that $b_t - l_t \geq 0$. Now, along the sequence generated by the adaptive forecasting and policy implementation process, it must be the case that $\lim_{t\to\infty} \Delta\left(X(t)\right) = 0$ (as the total social welfare gains is bounded which follows from the assumption that W is a continuous function over a compact set).

It follows that with probability one, $\lim_{t\to\infty} l_t = 0$. Therefore, $\sum_{t=0}^{\infty} l_t \leq \widehat{L}$, for some finite $\widehat{L}$, with probability one, across all the time periods when no mistake occurs.

Set $L = m_{\max}\widetilde{L} + \widehat{L}$. Then, gross social welfare losses are almost surely bounded by L.

E.2 Proof of Proposition 2

Let $X(t)$ be the current trajectory picked by the iterated application of growth diagnostics and let W_t be the associated level of social welfare. Let l_t be the social welfare loss incurred in period t, where if no social welfare loss occurs in period t then $l_t = 0$).

By assumption the total social welfare loss is almost surely bounded. Let the sequence V_t, indexed by t, be given by $V_t = W_t + \sum_{k=0}^{t} l_k$. Then this new sequence V_t is increasing. It is also bounded as it is the sum of two bounded sequences. Therefore it converges to a limit, say $\widehat{V}$. But this implies that W_t also converges to some limit $\widehat{W}$.

Let the associated sequence of trajectories be labelled as $\left(\widehat{X}(t) : t \geq 0\right)$ and let $\widehat{Y}$

denote any limit point of the sequence with $W(\widehat{Y}) = \widehat{W}$.

Under the assumption of almost surely bounded losses, there is a subsequence (with a slight abuse of notation, denoted in the same way as the original sequence) $\left(\widehat{X}(t) : t \geq 0\right)$ such that both $\lim_{t\to\infty} l_t = 0$ and $\lim_{t\to\infty} \Delta\left(\widehat{X}(t); d_t\right) = 0$.

Suppose, by contradiction, $W(\widehat{Y}) < W(Y)$. Then, starting from $\widehat{Y}$, it follows that $\Delta\left(\widehat{Y}; d_t\right) > 0$ at each $t \in T$. By continuity of W, there exists a T' such that for all $t > T'$, $\Delta\left(\widehat{X}(t); d_t\right) > 0$.

As losses are almost surely bounded, it follows that there exists T'' such that for all $t > T''$, $\Delta\left(\widehat{X}(t)\right) > 0$. But, then, $\widehat{Y}$ cannot be a limit point, a contradiction.

Therefore, $W(\widehat{Y}) = W(Y)$ and by continuity of $W(.)$ over $\mathbf{X}$, $W(\widehat{Y}) = \overline{W}$ and by the strict concavity of W, $\widehat{Y} = Y$.

Chapter 4

Conclusion

This thesis describes how smart information management can be utilised to improve outcomes for economic agents. Various mechanisms in diverse contexts are discussed to explore the potential to do so.

Chapter 1 studies a firm in which a manager may make systematic mistakes when prioritising adjustments to changing circumstances across different departments. I show theoretically that in this scenario the owner may improve profits by implementing the management technique 'management-by-exception', resulting in a restricted access to information for the manager. It should be noted that these results obtain without any inherent conflict of interests in the firm, but instead are mainly driven by psychological tendencies causing the manager to incorrectly evaluate information.

In Chapter 2, I discuss evidence from a computer laboratory experiment in the realm of household finance. In the real-world, private households tend to decrease their net profits from participating on the stock market by trading too much and, thus, incurring costs via commission fees, increased taxes and other positions. Participants in the experiment qualitatively emulate this behaviour in the laboratory despite being incentivised to maximise their gains from trading. I consider two information interventions to increase trading profits. First, I provide feedback, allowing subjects to more easily discover their own inability to increase profits from trading actively. Second, I reduce access to historic portfolio performance in order to curb overreactions to recent movements in stock value. Both interventions lead to significant changes in behaviour and increase net profits for participants.

Chapter 3 theoretically deals with a government trying to prioritise reforms. Despite being well-informed on any direct effects that a given policy-package may trigger, the government faces uncertainties on effects that may arise when changes in different realms of the economy interact. Hence, the government may actually make mistakes when implementing policies, causing overall social welfare to decrease with any given action. Even allowing for such mistakes, the chapter outlines conditions under which the policy-makers will still cause the economy to converge towards the social optimum. A key role in this process falls on the ability of policy-makers to use performance feedback on their policies in order to adjust the credibility of experts providing projections for future policy projects.

A large literature studies the propensity of controlling an agent's information channels towards guiding the agent's behaviour. Predominantly, these setups feature an inherent conflict of interests between the information designer and the receiver of the information. As a consequence, the information designer uses the control over the information channels to the detriment of the receiver. With this thesis, I hope to extend the scope of this method of persuasion. I show that purposefully manipulating an individual's inflow of information may mitigate negative consequences arising from systematic mistakes in the decision-making. It may thus help individuals making the correct choices for themselves and increase their outcomes without decreasing another individual's outcomes in return.

Bibliography

Alpert, Marc and Howard Raiffa, "A progress report on the training of probability assessors," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press, 1982, pp. 294–305.

Andries, Marianne and Valentin Haddad, "Information Aversion," *forthcoming in Journal of Political Economy*, 2019.

Antonakis, John, Bruce J Avolio, and Nagaraj Sivasubramaniam, "Context and leadership: an examination of the nine-factor full-range leadership theory using the Multifactor Leadership Questionnaire," *The Leadership Quarterly*, 2003, *14* (3), 261–295.

Athey, Susan, Joshua Gans, Scott Schaefer, and Scott Stern, "The Allocation of Decisions in Organizations," *Stanford University Research Paper No. 1322*, 1994.

Axelos, "The Axelos 2016 PRINCE2 Report," `https://www.axelos.com/Corporate/media/Files/AXELOS-PRINCE2-Report-2016.pdf` 2016. Accessed: 2018-07-19.

Barber, Brad M. and Terrance Odean, "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *The Journal of Finance*, 2000, *55* (2), 773–806.

_ and _ , "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," *The Quarterly Journal of Economics*, 2001, *116* (1), 261–292.

_ and _ , "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors," *The Review of Financial Studies*, 12 2008, *21* (2), 785–818.

_ , Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, "Just How Much Do Individual Investors Lose by Trading?," *The Review of Financial Studies*, 2008, *22* (2), 609–632.

Barber, Brad M, Yi-Tsung Lee, Yu-Jane Liu, Terrance Odean, and Ke Zhang, "Learning, Fast or Slow," *The Review of Asset Pricing Studies*, 08 2020, *10* (1), 61–93.

Bass, Bernard M., *Leadership and performance beyond expectations.*, New York: The Free Press, 1985.

_ , "From transactional to transformational leadership: Learning to share the vision," *Organizational Dynamics*, 1990, *18* (3), 19 – 31.

_ **and Bruce J. Avolio**, "Platoon Readiness as a Function of Leadership, Platoon, and Company Cultures," *United States Army Research Institute for the Behavioral and Social Sciences Technical Report 1104*, 2000.

Beggs, A. W., "Queues and Hierarchies," *The Review of Economic Studies*, 2001, *68* (2), 297–322.

Benartzi, Shlomo and Richard H. Thaler, "Myopic Loss Aversion and the Equity Premium Puzzle," *The Quarterly Journal of Economics*, 1995, *110* (1), 73–92.

Benmelech, Efraim and Carola Frydman, "Military CEOs," *Journal of Financial Economics*, 2015, *117* (1), 43 – 59.

Bergemann, Dirk and Stephen Morris, "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium," *The American Economic Review*, 2016, *106* (5), 586–591.

_ **and** _ , "Information Design: A Unified Perspective," *Journal of Economic Literature*, March 2019, *57* (1), 44–95.

Bergoeing, Raphael, Norman V. Loayza, and Facundo Piguillem, "The Whole is Greater than the Sum of Its Parts: Complementary Reforms to Address Microeconomic Distortions," *The World Bank Economic Review*, 2015, *30* (2), 268–305.

Bernile, Gennaro, Vineet Bhagwat, and P. Raghavendra Rau, "What Doesn't Kill You Will Only Make You More Risk-Loving: Early-Life Disasters and CEO Behavior," *The Journal of Finance*, 2017, *72* (1), 167–206.

Bertrand, Marianne and Antoinette Schoar, "Managing with Style: The Effect of Managers on Firm Policies," *The Quarterly Journal of Economics*, 2003, *118* (4), 1169–1208.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, "Does Aggregated Returns Disclosure Increase Portfolio Risk Taking?," *The Review of Financial Studies*, 2016, *30* (6), 1971–2005.

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sébastien Pouget, "Judgemental Overconfidence, Self-Monitoring, and Trading Performance in an Experimental Financial Market," *The Review of Economic Studies*, 2005, *72* (2), 287–312.

Blackwell, David, "Comparison of Experiments," in "Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability" University of California Press Berkeley, Calif. 1951, pp. 93–102.

_ , "An analog of the minimax theorem for vector payoffs," *Pacific Journal of Mathematics*, 1956, *6* (1), 1–8.

Blanco, Mariana, Patricio S. Dalton, and Juan F. Vargas, "Does the Unemployment Benefit Institution Affect the Productivity of Workers? Evidence from the Field," *Management Science*, 2017, *63* (11), 3691–3707.

Bloom, Nicholas and John Van Reenen, "Measuring and Explaining Management Practices across Firms and Countries," *The Quarterly Journal of Economics*, 2007, *122* (4), 1351–1408.

_ , Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen, "What Drives Differences in Management?," *mimeo*, 2017.

Bondt, Werner F. M. De and Richard Thaler, "Does the Stock Market Overreact?," *The Journal of Finance*, 1985, *40* (3), 793–805.

Bono, Joyce E. and Timothy A. Judge, "Personality and Transformational and Transactional Leadership: A Meta-Analysis," *Journal of Applied Psychology*, 2004, *89* (5), 901–910.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience Theory of Choice Under Risk," *The Quarterly Journal of Economics*, 2012, *127* (3), 1243.

_ , _ , and _ , "Salience and Consumer Choice," *Journal of Political Economy*, 2013, *121* (5), 803–843.

Brocas, Isabelle and Juan D. Carrillo, "Influence through ignorance," *The RAND Journal of Economics*, 2007, *38* (4), 931–947.

Burns, James MacGregor, *Leadership*, New York: Harper & Row, 1978.

Business Dictionary, "management by exception (MBE)," http://www.businessdictionary.com/definition/management-by-exception-MBE.html 2018. Accessed: 2018-07-19.

Calvó-Armengol, Antoni, Joan de Martí, and Andrea Prat, "Communication and influence," *Theoretical Economics*, 2015, *10* (2), 649–690.

Campos, Nauro F. and Abrizio Coricelli, "Growth in Transition: What We Know, What We Don't, and What We Should," *Journal of Economic Literature*, 2002, *40* (3), 793–836.

Cesa-Bianchi, N. and G. Lugosi, *Prediction, Learning, and Games*, Cambridge: Cambridge University Press, 2004.

Chen, An-Sing and Teng Yuan Cheng, "Are men more overconfident than women? Evidence from the Taiwan futures market," *mimeo*, 2018.

Cho, Seo-Young, "Explaining Gender Differences in Confidence and Overconfidence in Math," *mimeo*, 2017.

Copeland, Thomas E. and Daniel Friedman, "The Market Value of Information: Some Experimental Results," *The Journal of Business*, 1992, *65* (2), 241–266.

Cronqvist, Henrik and Stephan Siegel, "The genetics of investment biases," *Journal of Financial Economics*, 2014, *113* (2), 215 – 234.

Croson, Rachel and Uri Gneezy, "Gender Differences in Preferences," *Journal of Economic Literature*, June 2009, *47* (2), 448–74.

Dahlquist, Magnus, José Vicente Martinez, and Paul Söderlind, "Individual Investor Activity and Performance," *The Review of Financial Studies*, 2016, *30* (3), 866–899.

Danilov, Anastasia and Dirk Sliwka, "Can Contracts Signal Social Norms? Experimental Evidence," *Management Science*, 2017, *63* (2), 459–476.

de Clippel, Geoffroy, Kfir Eliaz, and Kareen Rozen, "The Silent Treatment," *mimeo*, 2017.

Deaves, Richard, Erik Lders, and Guo Ying Luo, "An Experimental Test of the Impact of Overconfidence and Gender on Trading Activity," *Review of Finance*, 09 2008, *13* (3), 555–575.

Dekker, S. W. A. and D. D. Woods, "To Intervene or not to Intervene: The Dilemma of Management by Exception," *Cognition, Technology & Work*, Sep 1999, *1* (2), 86–96.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, June 2009, *47* (2), 315–72.

Dessein, Wouter, "Authority and Communication in Organizations," *The Review of Economic Studies*, 2002, *69* (4), 811–838.

_ **and Andrea Prat**, "Attention in Organizations," in Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds., *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford Handbooks, 2016, pp. 675–697.

_ **and Tano Santos**, "Managerial Style and Attention," *mimeo*, 2016.

_ **, Andrea Galeotti, and Tano Santos**, "Rational Inattention and Organizational Focus," *American Economic Review*, June 2016, *106* (6), 1522–36.

Dittmar, Amy and Ran Duchin, "Looking in the Rearview Mirror: The Effect of Managers' Professional Experience on Corporate Financial Policy," *The Review of Financial Studies*, 2016, *29* (3), 565–602.

Dreze, J. H. and D. de la Vallee Poussin, "A Ttonnement Process for Public Goods," *The Review of Economic Studies*, 1971, *38* (2), 133–150.

Edwards, Sebastian, "The Sequencing of Economic Reform: Analytical Issues and Lessons from Latin American Experiences," *The World Economy*, 1990, *13* (1), 1–14.

Eichengreen, B. and M. Kawai, *Renminbi Internationalization: Achievements, Prospects, and Challenges*, Washington, DC: Brookings Institution Press, 2015.

Englmaier, Florian, Andreas Roider, and Uwe Sunde, "The Role of Communication of Performance Schemes: Evidence from a Field Experiment," *Management Science*, 2017, *63* (12), 4061–4080.

Fama, Eugene F., "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, 1970, *25* (2), 383–417.

Felipe, Jesus, Norio Usui, and Arnelyn Abdon, "Rethinking the Growth Diagnostics Approach: Questions from the Practitioners," *Journal of International Commerce, Economics and Policy*, 2011, *02* (02), 251–276.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Frydman, Cary and Baolian Wang, "The impact of salience on investor behavior: evidence from a natural experiment," *forthcoming in Journal of Finance*, 2019.

Gao, Xiaohui and Tse-Chun Lin, "Do Individual Investors Treat Trading as a Fun and Exciting Gambling Activity? Evidence from Repeated Natural Experiments," *The Review of Financial Studies*, 2014, *28* (7), 2128–2166.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, *108* (5), 874–904.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer, "Expectations and Investment," *NBER Macroeconomics Annual*, 2016, *30* (1), 379–431.

Gentzkow, Matthew and Emir Kamenica, "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, May 2016, *106* (5), 597–601.

Ghosal, Sayantan and James Porter, "Decentralised exchange, out-of-equilibrium dynamics and convergence to efficiency," *Mathematical Social Sciences*, 2013, *66* (1), 1 – 21.

_ and Massimo Morelli, "Retrading in market games," *Journal of Economic Theory*, 2004, *115* (1), 151–181.

Gillespie, Gary, "Inclusive Growth in Scotland," *mimeo, presentation at "Persistent Deprivation and Internal Constraints: Analysis and Policy Interventions", Workshop, Glasgow University, 12th February, 2016*, 2016.

Gneezy, Uri and Jan Potters, "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 1997, *112* (2), 631–645.

Grinblatt, Mark and Matti Keloharju, "Sensation Seeking, Overconfidence, and Trading Activity," *The Journal of Finance*, 2009, *64* (2), 549–578.

_ , _ , and Juhani T. Linnainmaa, "IQ, trading behavior, and performance," *Journal of Financial Economics*, 2012, *104* (2), 339 – 362. Special Issue on Investor Sentiment.

Hanke, Michael, Jürgen Huber, Michael Kirchler, and Matthias Sutter, "The economic consequences of a Tobin tax—An experimental analysis," *Journal of Economic Behavior & Organization*, 2010, *74* (1), 58–71.

Hausmann, R., B. Klinger, and R. Wagner, "Doing Growth Diagnostics in Practice: A 'Mindbook'," *CID Working Paper No. 177*, 2008.

_ , D. Rodrik, and A. Velasco, "Growth Diagnostics," *mimeo*, 2005.

Hausmann, Ricardo, Lant Pritchett, and Dani Rodrik, "Growth Accelerations," *Journal of Economic Growth*, 2005, *10* (4), 303–329.

Hayek, F. H., *The constitution of liberty*, Chicago: University of Chicago Press, 1978.

Heilmann, Sebastian, "From Local Experiments to National Policy: The Origins of China's Distinctive Policy Process," *The China Journal*, 2008, (59), 1–30.

_ , "Policy Experimentation in Chinas Economic Rise," *Studies in Comparative International Development*, 2008, *43* (1), 1 – 26.

Holt, Charles A. and Susan K. Laury, "Risk Aversion and Incentive Effects," *The American Economic Review*, 2002, *92* (5), 1644–1655.

Hottman, Stephen B. and Kari Sortland, "UAV Operators, Other Airspace Users, and Regulators: Critical Components of an Uninhabited System," in H. L. Pedersen, H. K. Connor, and E. Salas, eds., *Human Factors of Remotely Operated Vehicles*, New York: Elsevier Jai, 2006, pp. 71–88.

Huber, Jürgen, "J-shaped returns to timing advantage in access to information Experimental evidence and a tentative explanation," *Journal of Economic Dynamics and Control*, 2007, *31* (8), 2536–2572.

_ , Daniel Kleinlercher, and Michael Kirchler, "The impact of a financial transaction tax on stylized facts of price returns—Evidence from the lab," *Journal of Economic Dynamics and Control*, 2012, *36* (8), 1248–1266. Quantifying and Understanding Dysfunctions in Financial Markets.

_ , Michael Kirchler, and Matthias Sutter, "Is more information always better?: Experimental financial markets with cumulative information," *Journal of Economic Behavior & Organization*, 2008, *65* (1), 86–104.

Ivanov, Maxim, "Optimal Signals in Bayesian Persuasion Mechanisms," 2015.

Jowell, R., *Trying it out: the role of 'pilots' in policy-making*, London: Government Chief Social Researcher's Office, 2003.

Judge, Timothy A. and Ronald F. Piccolo, "Transformational and Transactional Leadership: A Meta-Analytic Test of Their Relative Validity," *Journal of Applied Psychology*, 2004, *89* (5), 755–68.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, October 2011, *101* (6), 2590–2615.

Karacaoglu, G., "The New Zealand Treasury's Living Standards Framework - Exploring a Stylised Model," *New Zealand Treasury Working Paper 15/12*, 2015.

_ , "Inclusive Growth in Scotland," *"Intergenerational Wellbeing and Public Policy: What Difference does Focusing on Wellbeing Make?", presentation at the "Policy Making and Wellbeing in New Zealand" Seminar, St. Andrews House, 9th December, 2016*, 2016.

Keynes, John M., *The General Theory of Employment, Interest and Money*, London, UK: Macmillan, 1936.

Kőszegi, Botond and Adam Szeidl, "A Model of Focusing in Economic Choice," *The Quarterly Journal of Economics*, 2013, *128* (1), 53–104.

Larson, Francis, John A. List, and Robert D. Metcalfe, "Can Myopic Loss Aversion Explain the Equity Premium Puzzle? Evidence from a Natural Field Experiment with Professional Traders," *mimeo*, 2016.

Li, Wei and Dennis Tao Yang, "The Great Leap Forward: Anatomy of a Central Planning Disaster," *Journal of Political Economy*, 2005, *113* (4), 840–877.

Lin, J. Y., F. Cai, and Z. Li, *The China Miracle: Development Strategy and Economic Reform*, Hong Kong: Chinese University Press, 2003.

Lindblom, Charles E., "The Science of "Muddling Through"," *Public Administration Review*, 1959, *19* (2), 79–88.

_ , "Still Muddling, Not Yet Through," *Public Administration Review*, 1979, *39* (6), 517–526.

Linnainmaa, Juhani T., Brian T. Melzer, and Previtero Alessandro, "The Misguided Beliefs of Financial Advisors," *mimeo*, 2018.

Lipsey, R. G. and Kelvin Lancaster, "The General Theory of Second Best," *The Review of Economic Studies*, 1956, *24* (1), 11–32.

List, J. and U. Gneezy, *The Why Axis: Hidden Motives and the Undiscovered Economics of Everyday Life*, London: Random House Books, 2014.

Liu, Dahai, Christopher Reynolds, Dennis Vincenzi, and Shawn Doherty, "Effect of Pilot and Air Traffic Control Experiences and Automation Management Strategies on Unmanned Aircraft Systems Mission Task Performance," *Human Factors and Ergonomics in Manufacturing & Service Industries*, 2013, *23* (5), 424–435.

Lowe, Kevin B., K. Galen Kroeck, and Nagaraj Sivasubramaniam, "Effectiveness correlates of transformational and transactional leadership: A meta-analytic review of the mlq literature," *The Leadership Quarterly*, 1996, *7* (3), 385–425.

Mackintosh, Donald P., *Management by exception*, Englewood Cliffs, NJ: Prentice-Hall, 1978.

Malinvaud, E., "Prices for Individual Consumption, Quantity Indicators for Collective Consumption," *The Review of Economic Studies*, 1972, *39* (4), 385–405.

Malmendier, Ulrike and Geoffrey Tate, "CEO Overconfidence and Corporate Investment," *The Journal of Finance*, 2005, *60* (6), 2661–2700.

_ , _ , and Jon Yan, "Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies," *The Journal of Finance*, 2011, *66* (5), 1687–1733.

March, James G. and Herbert Simon, *Organizations*, New York: John Wiley, 1958.

Martin, Ian W. R. and Robert S. Pindyck, "Averting Catastrophes: The Strange Economics of Scylla and Charybdis," *The American Economic Review*, 2015, *105* (10), 2947–2985.

Micklin, Philip and Nikolay V. Aladin, "Reclaiming the Aral Sea.," *Scientific American*, 2008, *298* (4), 64–71.

Micklin, Philip P., "Desiccation of the Aral Sea: A Water Management Disaster in the Soviet Union," *Science*, 1988, *241* (4870), 1170–1176.

Milgrom, Paul and John Roberts, "Relying on the Information of Interested Parties," *The RAND Journal of Economics*, 1986, *17* (1), 18–32.

Mukand, Sharun W. and Dani Rodrik, "In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance," *The American Economic Review*, 2005, *95* (1), 374–383.

Mullainathan, Sendhil and Eldar Shafir, *Scarcity: Why Having Too Little Means So Much*, New York, NY: Times Books, 2013.

National Bureau of Statistics of China, `http://data.stats.gov.cn/english/easyquery.htm?cn=C01` 2017. accessed 25th of January, 2017.

National Cotton Council of America, `http://www.cotton.org/econ/cropinfo/cropdata/rankings.cfm` 2017. accessed 25th of January, 2017.

Naughton, B., *The Chinese Economy: Transitions and Growth*, Cambridge: MIT Press, 2006.

North, D. C., *Institutions, institutional change and economic performance*, Cambridge: Cambridge University Press, 1990.

Ocasio, William, "Towards an attention-based view of the firm," *Strategic Management Journal*, 1997, *18* (S1), 187–206.

Odean, Terrance, "Volume, Volatility, Price, and Profit When All Traders Are Above Average," *The Journal of Finance*, 1998, *53* (6), 1887–1934.

_ , "Volume, Volatility, Price, and Profit When All Traders Are Above Average," *The Journal of Finance*, 1998, *53* (6), 1887–1934.

Palley, Thomas, "Speculation and Tobin taxes: Why sand in the wheels can increase economic efficiency," *Journal of Economics*, 1999, *69* (2), 113126.

Patterson, Coleman E. P., Jerry Bryan Fuller, Kim Hester, and Donna Y. Stringer, "A Meta-Analytic Examination of Leadership Style and Selected Follower Compliance Outcomes," *mimeo*, 1995.

Pigou, A.C., "Some Aspects of the Welfare State," *Diogenes*, 1954, *2* (7), 1–11.

Powell, Michael, "An Influence-Cost Model of Organizational Practices and Firm Boundaries," *The Journal of Law, Economics, and Organization*, 2015, *31* (S1), i104–i142.

Rantakari, Heikki, "Governing Adaptation," *The Review of Economic Studies*, 2008, *75* (4), 1257–1285.

Reis, Ricardo, "Inattentive Consumers," *NBER Working Paper 10883*, 2004.

Rodrik, Dani, "Second-Best Institutions," *The American Economic Review*, 2008, *98* (2), 100–104.

— , "Diagnostics before Prescription," *Journal of Economic Perspectives*, 2010, *24* (3), 33–44.

Roland, G., *Transition and economics: politics, markets, and firms*, Cambridge: MIT Press, 2000.

Sah, Raaj Kumar and Joseph E. Stiglitz, "The Architecture of Economic Systems: Hierarchies and Polyarchies," *The American Economic Review*, 1986, *76* (4), 716–727.

Scottish Government, "Inclusive Growth Diagnostic: Boosting Competitiveness and Employment Opportunities," *mimeo*, 2016.

Seru, Amit, Tyler Shumway, and Noah Stoffman, "Learning by Trading," *The Review of Financial Studies*, 09 2009, *23* (2), 705–739.

Simon, Herbert, *Administrative Behavior*, New York: The Free Press, 1945.

Sims, Christopher A., "Implications of rational inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665 – 690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Taylor, Shelley E. and Susan T. Fiske, "Salience, Attention, and Attribution: Top of the Head Phenomena," in Leonard Berkowitz, ed., *Advances in Experimental Social Psychology*, Vol. 11, Academic Press, 1978, pp. 249–288.

Techblog, Netflix, `http://techblog.netflix.com/2016/04/its-all-about-testing-netflix.html` 2017. accessed 26th of January, 2017.

Thaler, Richard H. and H. M. Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy*, 1981, *89* (2), 392–406.

Tobin, James, "A Proposal for International Monetary Reform," *Eastern Economic Journal*, 1978, *4* (3/4), 153–159.

Towne, Henry R., "The engineer as economist.," *Transactions of the American Society of Mechanical Engineers*, 1886, *7*, 428–441.

Wolfe, Jeremy M., Serena J. Butcher, Carol Lee, and Megan Hyle, "Changing Your Mind: On the Contributions of Top-Down and Bottom-Up Guidance in Visual Search for Feature Singletons," *Journal of Experimental Psychology: Human Perception and Performance*, 2003, *29* (2), 483–502.

World Bank, *Economic Growth in the 1990s: Learning from a Decade of Reform*, Washington, DC: World Bank, 2005.