Essays on the Economics of Information and Education

Claudia Herresthal

Department of Economics and University College,

University of Oxford

A thesis under the supervision of Dr. Margaret Meyer,

submitted in fulfilment of the degree of

Doctor of Philosophy in Economics[1]

Michaelmas Term 2016

[1]Approximate word count: 85,500 words.

Abstract

This thesis consists of three self-contained chapters which study different topics in the economics of information and education.

The first chapter studies what families can infer about schools' relative quality from rankings, when schools' performance depends not only on their quality but also on their student intake. I develop a dynamic framework in which each cohort of families has access to the most recent ranking, but can neither observe application choices nor the allocation of past students. I study a steady-state Bayesian-Nash equilibrium, and find that a performance-based ranking is more informative about school quality if oversubscribed schools select a larger share of their intake based on ability. I also find that such a ranking is more informative if it is less costly for families to attend a non-local school.

The second chapter studies a decision maker (DM) who can consult an advisor before choosing whether or not to switch from the status quo. Both agree that switching is optimal if and only if a certain hypothesis is true. But the advisor may be biased in how he trades off falsely accepting against falsely rejecting the hypothesis. The advisor has access to a testing technology which generates noisy but informative binary signals. Over two periods, the advisor can sequentially carry out costly tests before reporting his findings to the DM. I contrast the setting in which the DM observes all test outcomes with a setting in which testing itself is hidden and the advisor verifiably discloses outcomes. I fully characterise the circumstances in which the DM is strictly better off with hidden testing than with observable testing. Finally, I study what my findings imply about the bias of the ideal advisor.

The third chapter studies the influence of peers on educational performance. I estimate a linear-in-means model of peer effects using a methodology adopted from De Giorgi et al. [2010]. For identification, I exploit the fact that students have different classmates in different subjects. As instruments for classmates' performance I use the characteristics of students who share common classmates with the individual but who are not themselves the individual's classmates. I find that students are significantly influenced by their classmates' average performance but not by their classmates' average characteristics: a student's average test score in maths and English increases by 0.677 of a standard deviation if the average performance of his peer group increases by one standard deviation.

Acknowledgements

First and foremost, I want to thank my supervisor, Dr. Meg Meyer, for her invaluable guidance and advice throughout my MPhil and DPhil studies. Her extraordinary ability to understand models quickly and her immense knowledge of the literature continue to fascinate me and are qualities I aspire to develop. I also greatly appreciate her support with my job search, which has led to the opportunity to continue doing research beyond the DPhil.

I am grateful to Professor Simon Burgess for providing me with an opportunity to conduct empirical work, in particular, for sharing the data set he helped collect and for the interesting discussions about the identification of peer effects, which provided the inspiration for the third chapter.

I would like to thank my fellow students for their companionship and the insights they contributed. I am also grateful for helpful comments and feedback from other members of the department, visiting speakers and conference audiences. For many discussions, feedback on my work and advice on how to navigate academia, I am especially grateful to Anna Sanktjohanser, Collin Raymond, Inga Deimen, Pawel Dziewulski, Ole Jann, Nikki Shure, Pawel Gola, Abi Adams, Alex Teytelboym, James Wisson, Sam Jindani, Kristofer Hammerbäck, Felix Pretis and Nicholas Chesterley.

I gratefully acknowledge funding from the Department of Economics, University College and the Royal Economic Society, which made my DPhil work possible. I am also grateful to Queens' College, Cambridge, for providing me with an excellent working environment for the final stretch.

The support of friends and family in Oxford and beyond helped make this DPhil a reality. I am grateful to "the nest" and its associate members. I also want to express my special gratitude towards Lucy Ambrose, Teresa Kykre-Smith, Patrick Watts and Lilja Kristinsdottir. Finally, my greatest thanks go to my mum and dad and my sister for their endless love and support.

Contents

1 Performance-based Rankings and School Quality **3**

1.1 Introduction 4

1.2 Model 9

1.3 Illustrative Example 13

1.4 Optimal Application Strategy 16

1.5 Equilibrium Steady State 17

1.5.1 Convergence to Equilibrium Steady State 21

1.6 Comparative statics 23

1.7 Welfare 27

1.8 Externalities 30

1.8.1 Complements or Substitutes 31

1.9 Extensions and Limitations 32

1.9.1 Longer Window of Performance-based Rankings 32

1.9.2 No Ability-based Selection by Schools 33

1.9.3 Peer Effects 34

1.9.4 Signal of School Quality 35

1.9.5 Absorbing States 36

1.9.6 Replacement of Unpopular Schools 37

1.10 Conclusion 37

1.11 Appendix 38

1.11.1 Equilibrium Steady State: Proposition 1 38

1.11.2 Convergence to Equilibrium Steady State: Proposition 2 42

1.11.3 Comparative Statics: Proposition 3 . . . 46
1.11.4 Comparative Statics: Proposition 4 . . . 49
1.11.5 Welfare: Proposition 5 . . . 49
1.11.6 Welfare: Proposition 6 . . . 51
1.11.7 Extensions: Longer Window of Performance-based Rankings . . . 52
1.11.8 Extensions: No Ability-based Selection by Schools . . . 55
1.11.9 Extensions: Peer Effects . . . 56
1.11.10 Extensions: Signal of School Quality . . . 58

2 Hidden Testing and Selective Disclosure of Evidence 60

2.1 Introduction . . . 61
2.2 Model . . . 65
2.2.1 General Framework . . . 66
2.2.2 Observable Testing . . . 67
2.2.3 Hidden Testing . . . 68
2.3 No Conflict of Interest - Single Player Benchmark . . . 69
2.4 Conflict of Interest with Observable Testing . . . 74
2.5 Conflict of Interest with Hidden Testing . . . 82
2.6 Choosing an Advisor . . . 94
2.7 Related Literature . . . 97
2.7.1 Non-sequential Testing . . . 97
2.7.2 Sequential Testing . . . 100
2.8 Conclusion . . . 102
2.9 Appendix . . . 104
2.9.1 Lemma 1 . . . 104
2.9.2 Lemma 2 . . . 113
2.9.3 Lemma 3 . . . 116
2.9.4 Theorem 1 . . . 128
2.9.5 Proposition 1 . . . 130
2.9.6 Proposition 2 . . . 130
2.9.7 Proposition 3 . . . 130

2.9.8 Proposition 4 . . . 131

3 Peer Effects in GCSE Performance **133**

3.1 Introduction . . . 134

3.2 Identification Problems and Strategies . . . 137

3.2.1 Identification Problems . . . 137

3.2.2 Exploiting Partially Overlapping Peer Groups . . . 138

3.3 Data . . . 142

3.4 Peers . . . 144

3.5 Empirical Strategy . . . 152

3.6 Results . . . 155

3.6.1 Outcome measure: Average GCSE Test Score in English, Maths and Science . . . 155

3.6.2 Outcome measure: Subject-specific GCSE Test Scores . . . 158

3.6.3 Discussion . . . 162

3.7 Robustness . . . 164

3.7.1 Standard Errors Clustered at the Science Class Level . . . 164

3.7.2 Equally Weighted Peers . . . 164

3.7.3 Gender as a Determinant of Class Allocation . . . 165

3.8 Conclusion . . . 165

3.9 Appendix . . . 170

3.9.1 Descriptive Statistics of Peer and Excluded Peer Average Variables . . . 170

3.9.2 First-Stage Regression for Table 3.12 . . . 172

3.9.3 LIML Estimation - Clustered Standard Errors . . . 173

3.9.4 Equally Weighted Peers - No Clustering . . . 174

3.9.5 Gender as a Determinant of Class Allocation - No Clustering . . . 175

Introduction

This thesis is a compilation of three essays which study different topics in the economics of information and education. Each essay represents a self-contained paper. Nevertheless, there are overarching aspects that establish links in between these papers. The first paper addresses the themes of information and education. It theoretically analyses how families choose between schools based on school rankings in a dynamic framework. The second paper also theoretically studies decision making in a dynamic framework under incomplete information. In particular, it addresses the question of how much transparency is optimal when an advisor sequentially gathers evidence for an uninformed decision maker, e.g. as in the context of scientific research. The third paper deals with peer effects in schools and thus belongs to the field of education. It is linked to the first because it investigates another determinant of educational achievement, but unlike the other two papers it uses an empirical approach.

The paper "Performance-based Rankings and School Quality" studies what families can infer about schools' relative quality from rankings, when schools' performance also depends on their student intake. I develop a dynamic framework in which each cohort of families has access to the most recent ranking, but can neither observe application choices nor the allocation of students in the past. The analysis takes into account that families' inferences will have an impact on the allocation of students and hence on future rankings, thereby affecting the inference problem encountered by future families. I find that, in the long run, a performance-based ranking of schools is more informative about quality if oversubscribed schools select a larger share of their intake based on ability. In addition, I also find that such a ranking is more informative if it is less costly for families to attend a non-local school. In further analysis, I allow families' demand to have an impact on the quality of schools and analyse welfare effects. I find that, if oversubscribed schools select a larger share of students based on their ability, the average quality of schools improves, but low-ability students are worse off.

The paper "Hidden Testing and Selective Disclosure of Evidence" compares how different levels of

transparency affect to what extent an advisor (he) can strategically manipulate a decision maker (she). The decision maker (DM) is uncertain whether or not to change her action and can consult an advisor to gather evidence. Both agree that changing action is optimal if and only if a certain hypothesis is true. But the advisor may be biased in how he trades off falsely accepting against falsely rejecting the hypothesis. Over the course of two periods, the advisor can sequentially acquire noisy but informative signals, each at a fixed cost. I contrast the setting in which the DM observes all test outcomes with a setting in which testing itself is hidden and the advisor verifiably discloses outcomes. I find that the DM may be strictly better off if the advisor gathers evidence in private and I fully characterise the circumstances in which this is the case. Depending on how strongly the DM is predisposed to stick with the status quo, hidden testing can be beneficial either if the advisor is more enthusiastic or more reluctant to change action than the DM. Finally, I study what my findings imply for the bias of the ideal advisor.

The paper "Peer Effects in GCSE Performance" estimates how much a student's peers influence his educational performance. Peer effects are challenging to identify. A student influences his peers and his peers in turn influence him, making it difficult to distinguish between the cause and the effect of peers' outcomes on the student's outcome. I estimate a linear-in-means model of peer effects using a methodology adopted from De Giorgi et al. [2010]. For identification, I exploit the fact that students have different classmates in different subjects. As instruments for classmates' performance I use the characteristics of students who share common classmates with the individual but who are not themselves the individual's classmates. I find that students are significantly influenced by their classmates' average performance but not by their classmates' average characteristics: a student's average test score in maths and English increases by 0.677 of a standard deviation if the average performance of his peer group increases by one standard deviation.

This thesis is structured as follows: the first chapter is the paper "Performance-based Rankings and School Quality", the second chapter is the paper "Hidden Testing and Selective Disclosure of Evidence" and the third chapter is the paper "Peer Effects in GCSE Performance" .

Chapter 1

Performance-based Rankings and School Quality

Abstract

Families can express a preference as to the school their child will attend and school rankings are published to help families learn about the quality of schools. However, school rankings are based on past performance, and past performance not only depends on school quality but also on the ability of the past intake. I study a dynamic framework in which each cohort of families has access to the most recent ranking, but can neither observe application choices nor the allocation of students in the past. The analysis takes into account that families' inferences about school quality will have an impact on student allocations and hence on future rankings, thereby affecting the inference problem encountered by future families. I study a steady-state Bayesian-Nash equilibrium, and find that a performance-based ranking is more informative about school quality if oversubscribed schools select a larger share of their intake based on ability. I also find that such a ranking is more informative if it is less costly for families to attend a non-local school. When I allow for supply side responses, I find that if oversubscribed schools select a larger share of students based on their ability then the average quality of schools improves, but low-ability students are worse off.

1.1 Introduction

Recent reforms allow families to send their child to a non-local school and have led to the introduction of performance-based school rankings.[1] These rankings are intended to inform families about schools' intrinsic quality, because it is important that families can compare schools in terms of the quality of their teaching and management. However, schools' performance is also influenced by the ability of their student intake. This makes it more difficult for families to infer school quality from rankings, because families usually observe neither the allocation nor the application choices of past students, and may not even observe the ranking based on which past students chose between schools.[2]

Given these informational constraints, families' inferences will depend on schools' admission policies. If oversubscribed schools were to allocate places by lottery rather than select students based on ability, families would expect schools to have more similar student intakes. Can this help them to better identify high-quality schools based on a school ranking? In addition, families' inferences will depend on how important school quality was for families in the past. When choosing a school, families take into account not only the school's quality but also other features which distinguish schools, such as distance to home or fields of specialisation. For example, if families in the past had a strong preference to attend their local school irrespective of its rank, families today would also expect schools to have more similar intakes and hence they would interpret rankings accordingly.

It follows that families' inferences can be affected by policymakers via similar channels. In particular, policymakers can regulate school admissions, e.g. they could require schools to give priority to less able local students over more able non-local students. In addition, they can affect how attractive it is for families to select schools based on indications of quality rather than other features, e.g. by subsidising transport costs or by reducing the bureaucratic burden associated with attending a non-local school.

This paper develops a dynamic framework which permits to keep track of what families infer about school quality from rankings, while also capturing that their inferences will have an impact on the student allocation and hence on future rankings, thereby affecting the inference problem encountered by future families. Given this framework, I determine the likelihood with which a higher-ranked school is of better

[1] E.g. as a result of expanding families' choice of secondary schools, 15% of students in the US do not attend a district-assigned school, and over half of students in England do not attend their nearest state school (Greene et al. [2010], Burgess et al. [2006]). In the US, school report cards were first required by the No Child Left Behind Act in 2002, and in England, school rankings have been released by the media annually since 1992.

[2] To help families, estimates of schools' value-added have been published. However, these estimates are rarely used by families (Coldron et al. [2008], Imberman and Lovenheim [2013]), and although they may reduce families' inference problem they do not completely eliminate it (Kane et al. [2002], Wilson and Piebalga [2008]).

quality in the limit of this dynamic process. I then analyse how this measure of the informativeness of a school ranking is influenced by changes in i) the proportion of an oversubscribed school's intake selected based on ability and ii) by the distribution of transport costs, where transport costs measure the additional costs incurred by a family if their child were to attend their non-local rather than local school.

A key insight is that the informativeness of rankings increases the more places at an oversubscribed school are allocated based on ability, e.g. based on an entrance exam. It may seem counterintuitive at first. After all, if places were allocated by lottery, schools would have similar intakes on average, and better performance could only be attributed to better school quality or to chance. By contrast, if an oversubscribed school got to select more able students, then student intakes would vary more across schools and better performance could also be attributed to better students. This suggests that selection on ability would make it harder for families to draw inferences about schools' relative quality. However, this reasoning breaks down when I consider a dynamic setting. If in each cohort an oversubscribed school selects the most able applicants then a school with strong recent performance is likely to have faced high demand in the past. This indicates that families in the past expected the school to be of higher quality and their expectations are on average correct. As families believe that better schools are likely to be attended by more able students, families are more confident that the relative performance of schools reflects their relative quality.

In the model of this paper, there are two districts, each containing one school of either good or bad quality which serves successive cohorts of families. Each cohort consists of a continuum of families, a fixed share of which has high-ability and the remainder of which has low-ability children. Each school has capacity for half the students. Students enrol where they apply unless the school is oversubscribed. In this case, successful applicants are selected according to the school's admission policy and rejected applicants attend the other school. A school's rank depends on its relative intrinsic quality and on how the ability distribution of its intake differs from that of the other school. Families perceive schools as horizontally differentiated, which is captured by the assumption that families are distributed along a Hotelling line and incur transport costs based on their distance from the school their child attends. Families prefer a good to a bad school, but dislike transport costs, and their preferences are distributed independently of their child's ability.[3] A priori, families are uninformed about the quality of each school, but they observe which school performed better with its previous intake. Hence, families may face a trade-off between

[3]This trade-off between quality and distance is well-documented, e.g. see Burgess et al. [2015]. Transport costs can also represent differences in families' valuation of other attributes, e.g. some value sports grounds, others fine art supplies.

a higher-ranked but more distant school and a lower-ranked local school. Crucially, families base their inferences about school quality only on the most recent ranking, and they do not observe the ability of previous student intakes or their application choices.

In equilibrium, the higher-ranked school receives applications from all local families and, in addition, from some proportion of non-local families. The proportion of non-local families who apply to the higher-ranked school will be referred to as the level of mobility and this level of mobility fully characterises application strategies of this cohort. Over time, the level of mobility increases and converges in the limit. At this point, families' application strategies are optimal given how informative they believe the ranking to be and their beliefs are consistent given the application strategies of families in the past. Instead of deriving the limit of this process, I derive an equilibrium steady-state in which these properties are also satisfied. In particular, I consider a steady state defined by the feature that the level of mobility is constant across cohorts. If families conjecture to be in such a steady state, then each cohort will find the ranking equally informative. To solve for an equilibrium steady-state, I find a fixed point of mobility such that this level of mobility characterises families' optimal application strategies given their beliefs and their beliefs are consistent with being in a steady state defined by this level of mobility.

I find that, in equilibrium steady-state, families are weakly better informed about the relative quality of schools if families perceive schools to be more similar in features unrelated to quality, i.e. if the distribution of transport costs would shift down (in the sense of first-order stochastic dominance). Therefore, reducing the barriers to choice and making it easier for families to pick the school of highest expected quality facilitates learning about school quality over time. In addition, I find that families are better informed about the relative quality of schools if oversubscribed schools select a higher proportion of their intake based on ability. This suggests that admission policies such as giving priorities to local students or allocating places by lottery may have negative implications for how informative rankings are in the long run.

The analysis also reveals an informational externality across different cohorts. The more families sought out the better-performing school in the past, the more informative is the performance-based ranking of schools as a signal of schools' relative quality for the current cohort.

The reforms which triggered the introduction of school rankings were intended to put pressure on low-quality schools to improve such that the overall quality of schools would increase. So far, I have focused on analysing to what extent demand will be sensitive to quality, holding the qualities of schools

fixed. In further analysis, I incorporate supply side effects by assuming that an undersubscribed school is replaced by a new school with some probability and the new school's quality is drawn from the same distribution as the original school's quality. Therefore, I do not simply assume that schools improve as a response to being undersubscribed. Instead, the overall quality of schools improves if families are more likely to identify the relatively better school, because then the school being replaced is more likely to be a low-quality than a high-quality school. The replacement of a school is not necessarily meant to represent a school closing down, but could also capture a change in quality due to new teachers or a new principal.

I find that the average quality of schools improves and both high- and low-ability students are better off when an undersubscribed school is replaced with some positive probability, compared to the case in which demand has no impact on schools' qualities. If i) an oversubscribed school selects a larger share of their intake based on ability or ii) if transport costs are lower, the average quality of schools improves but the benefits accrue only to high-ability students. Hence, policymakers may face a trade-off: allowing selection by schools improves how informative rankings are and puts pressure on worse schools to improve, however, the benefits of better school quality then accrue disproportionately to more able students, because they are more likely to gain access to better schools.

While families' choice of school is my leading example, the insights developed in this paper also apply to the context of patients' choice of hospital. Patients have to infer the quality of treatment from outcome measures of medical procedures, such as mortality rates, and these outcomes also depend on the health of patients treated in the past.

The predictions of this paper closely match empirical findings in health care choice by Chandra et al. [2016], who find that i) there is a correlation between a hospital's market share and its quality at a given point in time and that ii) higher-quality hospitals grow more over time. Further, they find that both these results are stronger when excluding emergency admissions. Finally, they propose that the publication of hospital report cards may have triggered these effects. In this paper, I provide a theoretical foundation for their claim. I explicitly model the inference made by patients from rankings of hospitals and use mobility to measure the expected correlation between market share and quality. I predict that positive mobility is associated with a higher average quality of hospitals. In addition, when hospitals are perceived to be less horizontally differentiated (lower transport costs), as is the case for non-emergency relative to emergency admissions, I predict this effect to be even stronger.

This paper contributes to the literature studying how families choose between schools, or how patients

choose between hospitals. Empirical research has estimated how sensitive consumers' demand is to past performance,[4] but has not considered that what consumers infer about quality from past performance may vary depending on how consumers are matched with providers. The existing theoretical literature has not yet studied families' inference in a dynamic setting, and in particular, it has not yet studied how the choice of school made by families today imposes an informational externality on the inference made by future families. Gavazza and Lizzeri [2007] study theoretically how transparency about school performance may affect schools differently depending on whether they can select between students, but they do not consider how the allocation of students to schools affects the informativeness of rankings over time. De Fraja and Landeras [2006] study effects on attainment when families choose between schools based on rankings, but they do not incorporate how the allocation between students and schools affects what families infer about teaching quality from rankings. Although this paper does not form part of the literature on matching algorithms,[5] it adds the important insight that the allocation of families to schools influences future families' beliefs about schools' relative quality and hence what preferences they will submit to the algorithm in the future.

The paper also improves our understanding of how market-based accountability reforms affect the quality of schools available to different families, when families are incompletely informed about school quality. Despite numerous empirical studies,[6] there has been little theoretical investigation of this issue. An exception is Barseghyan et al. [2014], who analyse school choice under the assumption that the quality of the school is defined by the quality of its students. By contrast, this paper assumes families care about the quality of teaching and analyses how the allocation of past students, who are not peers, affects families' inference about quality. An extension considers a model with peer effects.

The model has some similarities to models in the social learning literature,[7] since families base their decision on rankings which are an endogenous signal of the decisions made by earlier cohorts. However, unlike in the social learning literature, families do not observe their predecessors' choices directly. In addition, families also do not receive an exogenous private signal about the quality of schools, instead, they have private information because each cohort only observes the most recent ranking. Lobel et al. [2007] also model agents with a limited window of observation, but their focus is on conditions for

[4] E.g. for education see Black [1999], Bayer and McMillan [2005], Burgess et al. [2015], Imberman and Lovenheim [2013], Hastings and Weinstein [2007], for health care see Gaynor et al. [2012]

[5] E.g. Abdulkadiroglu and Sönmez [2003]

[6] E.g for education see Belfield and Levin [2002], Hastings and Weinstein [2007], Gibbons et al. [2006], for health care see Dranove et al. [2003], Cutler et al. [2004], Kessler and McClellan [2000], Gaynor et al. [2013]

[7] E.g. see Bikhchandani et al. [1992].

convergence whereas my focus is on the steady-state analysis. Callander and Hörner [2009] propose a steady-state analysis but focus on agents inferring information from the relative frequency with which actions were taken by predecessors. Formally, the results about learning in this paper are related to the results derived from the study of biased contests (Meyer [1991]), but, in this model, the bias is not directly set by a planner but arises endogenously in equilibrium.

This paper starts with an outline of the model in Section 1.2. It first introduces the model families' inference problem in an illustrative example in Section 1.3 before solving for equilibrium and conducting comparative statics and welfare analysis in Sections 1.4-1.7. Externalities are illustrated in section 1.8 by analysing how a social planner would assign families if faced with the same informational constraints. Finally, robustness checks can be found in Section 1.9 and analyse situations in which families observe a longer window of observations, in which they care about peers, and in which schools cannot select among applicants based on ability. In addition, I discuss in how far I can generalise assumptions on how schools' performance is determined, and on how school qualities are distributed. Section 1.10 concludes. All proofs can be found in the Appendix in Section 1.11.

1.2 Model

Successive generations (cohorts) of families are matched with schools. In each generation, each family applies to one of two schools. An oversubscribed school chooses whom to accept. Each family's payoff depends on both the quality and proximity of the school to which their child is admitted.

A general model will be presented in which an undersubscribed school is replaced by a new school with some probability $\gamma \geq 0$. The idea is that a school which cannot attract enough applications is under pressure to change its administration or teacher body. The general model nests the baseline model in which I study families' demand when school qualities are exogenously determined, i.e. $\gamma = 0$.

Schools **-** There are two (non-strategic) schools which differ in quality and location. The location of each school is fixed over time. Each school's quality is either good (G) or bad (B). Initially, each school's quality is drawn independently at random such that $Pr(G) = \theta$. If $\gamma > 0$, then when a school is replaced, a new school opens (in the same location) and its quality is also drawn independently at random such that $Pr(G) = \theta$. If oversubscribed, a school selects among applicants according to an exogenously given rule.

Families - Families are risk-neutral Bayesian decision makers. In each generation, there is a unit mass of families with high-ability students (H) and a unit mass of families with low-ability students (L). Families' preferences over schools are distributed independently of students' abilities. Families perceive schools as both horizontally and vertically differentiated. Schools are perceived as vertically differentiated because each family receives a benefit $V > 0$ if their student attends a good school and 0 otherwise. Further, schools are perceived as horizontally differentiated because families are distributed along a Hotelling line and the costs of attending a school depend on their location. In particular, the additional costs associated with attending the school that is further away (non-local school) rather than the closer school (local school) are captured by the level of transport costs $c \geq 0$ and are distributed among families according to a continuous distribution function $F(c)$. The distribution of costs of attending the non-local rather than the local school is symmetric with respect to the identity of the local school. The payoff of a family with transport costs c varies with the characteristics of the school attended as follows:

		Quality	
		good	bad
Proximity	local	V	0
	non-local	$V-c$	$-c$

Timing and Information - Generation 0 is the first generation of families to apply to schools. In generations $t \geq 0$, families do not observe schools' qualities, but each family observes their level of transport costs. Further, they know the distribution from which school qualities are drawn, as well as the distributions of student abilities and transport costs. In generation $t \geq 1$, each generation t observes schools' relative performance with students in generation $t-1$, and knows how schools' relative performance is stochastically dependent on schools' qualities and the ability of their student intakes. The school which ranked higher based on their performance in generation $t-1$ is referred to as the period-$t-1$ winner and denoted by W_{t-1}. Families do not observe past generations' application choices or past allocations of students to schools.[8]

Each family applies to one school. Schools have an equal capacity of a unit mass. All applicants are accepted unless applications exceed capacity. In this case, the oversubscribed school admits all local high-ability applicants, a proportion p out of the pool of non-local high-ability applicants and for the

[8]The most recent ranking is the most relevant and most easily accessible ranking for families, e.g. in England this ranking is circulated in the media. In the extensions, I analyse the case in which families consider a longer window of past rankings.

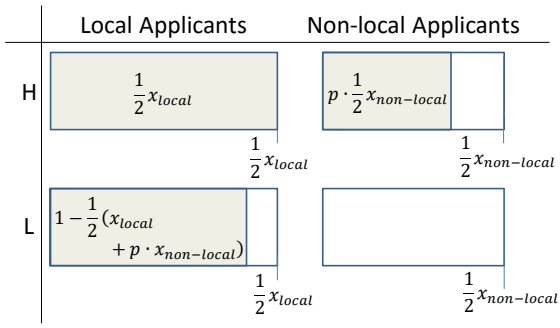

Figure 1.1: School Admissions - The boxes illustrate the different categories of applicants to the oversubscribed school, where x_{local} ($x_{non-local}$) denotes the mass of local (non-local) students applying, and the shaded areas present the mass of students admitted.

remaining places gives priority to local over non-local applicants independent of ability (see Figure 1.1). Rejected applicants attend the other school. Parameter p measures the extent to which oversubscribed schools select on ability, e.g. $p = 0$ would correspond to no selection on ability.[9]

After applications are received, but before students are admitted to schools, the undersubscribed school is replaced by a new school with probability $\gamma \in [0,1]$, whose quality is again drawn independently at random with $Pr(G) = \theta$.[10] Then performance and payoffs are realised and a school ranking (W_t) is constructed. Families know how the replacement of schools is determined, but they do not observe whether or not a replacement has taken place.

I will refer to the case in which $\gamma = 0$ as the case in which school quality is *exogenously* determined, whereas I will refer to the case in which $\gamma > 0$ as the case in which school quality is *endogenously* determined. School quality is endogenous if $\gamma > 0$ because the likelihood that a school is replaced depends on families' application choices in equilibrium.[11] Raising parameter γ to a positive value therefore links the chances that a school will stay in the market to families' demand for places at this school.

Each generation of families corresponds to one period.

Production function - Which one of the two schools ranks higher (wins) depends on schools' relative

[9]For the results on learning about school quality, the important feature of the admission process is that an oversubscribed school takes on a larger share of high-ability students than the other school. Whether some local high-ability students are rejected, or some non-local low-ability students are accepted only alters how much transport costs have been incurred relative to how accurately school quality is being inferred.

[10]Whether replacement of a school occurs before or after students are admitted to schools is irrelevant in steady state.

[11]Formally, if $\gamma = 0$ then school qualities are fixed over time. However, in steady state, this is equivalent to an *exogenously* chosen school being replaced by a new school.

quality, the allocation of student abilities and residual noise.[12] First, consider the probability with which a good school wins against a bad school, denoted by $Pr\left(\overline{G}B\right)$. Denote by h the share of high-ability students attending the good school, which will be determined in equilibrium. Then

$$Pr\left(\overline{G}B\right) = \frac{s}{2} + \left(h - \frac{1}{2}\right) d$$

where $s \in (1,2)$ reflects the performance advantage due to superior quality and $d \in (0, 2-s)$ reflects the performance advantage due to a larger share of high-ability students. If the allocation of students is balanced across schools ($h = \frac{1}{2}$) then which school wins depends only on the relative quality of schools and residual noise. If the good school has admitted a larger share of high-ability applicants than the bad school, its chances of winning are biased upwards, whereas if it has admitted a lower share of high-ability students than the bad school its chances are biased downwards. A rise in s reflects that which school wins depends more strongly on schools' relative qualities rather than on student ability or residual noise. If it were true that $s = 1$ schools would effectively be of the same quality, whereas if it were true that $s = 2$ the good school would win for certain. Further, a rise in d reflects that which school wins depends more strongly on how high-ability students are distributed across schools rather than on schools' relative quality or residual noise. If it were true that $d = 0$ the allocation of students would have no impact on relative performance, whereas if it were true that $s = 1$ and $d = 1$ then a school which admits all high-ability students would be guaranteed to win. The analysis rules out these extreme cases in order for both schools' relative quality and the allocation of students to have an impact on relative performance. Since these influences on performance enter the probability of winning additively, the increase in the chance of winning derived from accepting more high-ability students is independent of the school's quality.[13]

Next, consider the case in which both schools are of the same quality.[14] I will assume that the chances of winning are independent of whether both schools are good or both are bad, and that in these cases neither school has an advantage due to superior quality, i.e. $s = 1$. Then the probability that the school with an endogenously determined share h of high-ability students wins is given by:

$$Pr\left(\overline{G}G\right) = Pr\left(\overline{B}B\right) = \frac{1}{2} + \left(h - \frac{1}{2}\right) d$$

[12] For evidence that there is residual noise in rankings see Kane et al. [2002].

[13] For further discussion of the functional form of the signal, see the discussion in Section 1.9.4.

[14] In steady state, it is irrelevant how the relative performance of schools of the same quality is determined, due to the symmetry of the set-up.

Families' strategies - Each family decides which school to apply to. Strategies in generation t can be characterised by the share of families who apply to the most recent winner, i.e. W_{t-1}, in the population of families who consider W_{t-1} to be their non-local school. I will refer to this share of families as the level of mobility in generation t, denoted by $m_t \in [0,1]$. Families' optimal strategies in generation t (and hence mobility m_t) will depend on their beliefs about how likely the most recent winner is a better rather than a worse school. The reason why strategies can be characterised by mobility and how they depend on beliefs about the informativeness of the ranking will be explained in detail in Section 1.4.

Equilibrium Concept - The equilibrium concept is a Bayesian-Nash equilibrium steady state, i.e. a steady state defined by the level of mobility being unchanged across cohorts, in which families' application choices are optimal given beliefs and beliefs are consistent with equilibrium steady state play. The construction of such an equilibrium steady state is explained in detail in Section 1.5.

1.3 Illustrative Example

The following example will illustrate how successive generations of families draw inferences about schools' relative quality from the performance-based ranking they observe, starting with the first generation with access to a ranking. In particular, I will compare how well informed the first generation with access to rankings is, generation 1, compared to the subsequent generation, generation 2. I will show that generation 2 is better informed than generation 1, despite the fact that each generation only has access to the most recent ranking. This example is useful to build intuition for how families draw inferences about school qualities in steady state, as analysed in detail in Section 1.5. In addition, it will highlight why analysing an equilibrium steady state makes comparative statics more tractable. To this end, the example will use a simplified set-up.

There is one good and one bad school in the market, and their qualities are fixed over time. Families know that one school is better, and have symmetric prior beliefs about which school this is. Further, there are two students in each generation, one high-ability student and one low-ability student. Each school has capacity for one student. If two students apply then the high-ability student is accepted with probability $\frac{1}{2}+p$, and the rejected student attends the other school. If the high-ability student attends the good school, the good school wins with probability $\frac{1}{2}(s+d)$, otherwise the good school wins with

probability $\frac{1}{2}(s-d)$. Assume there are no transport costs.[15]

Generation 0 has no access to rankings. Due to their symmetric prior beliefs about which school is good, each family picks a school at random to which they apply. Consequently, the high-ability student is equally likely to attend the good or the bad school.

Generation 1 observes a ranking of schools based on performance with students in generation 0. Further, generation 1 knows how likely it is that the good school ranks higher conditional on the ability of its intake in generation 0. To interpret the ranking, generation 1 conjectures that the high-ability student in generation 0 was equally likely to attend the good or the bad school, i.e.

$$Pr\left(H_{t=0}^{G}\right)=\frac{1}{2}$$

where $Pr\left(H_{t}^{G}\right)$ denotes the event that the high-ability student in generation t attended the good school. Based on this conjecture, generation 1 constructs the probability that the good school is the winner with students in generation 0, denoted by $Pr\left(\overline{G}B_{t=0}\right)$, as follows:

$$\begin{aligned} Pr\left(\overline{G}B_{t=0}\right) &= Pr\left(H_{t=0}^{G}\right)\left[\frac{1}{2}(s+d)\right]+\left[1-Pr\left(H_{t=0}^{G}\right)\right]\left[\frac{1}{2}(s-d)\right] \\ &= \frac{1}{2}\left[\frac{1}{2}(s+d)\right]+\frac{1}{2}\left[\frac{1}{2}(s-d)\right]=\frac{s}{2}\geq\frac{1}{2} \end{aligned}$$

As families have a symmetric prior about which school is good, $Pr\left(\overline{G}B_{t=0}\right)$ constitutes families' posterior belief that the most recent winner is the good school. Since the most recent winner is more likely to be good than bad, both students will apply to this school.

Generation 2 observes a ranking based on performance with students in generation 1, but not the earlier ranking based on performance with students in generation 0. To interpret the ranking, they form a belief about how likely it is that the high-ability student in generation 1 attended the good school. First, they conjecture that in generation 0 families chose between schools at random. Second, they conjecture that in generation 1 both students applied to the period-0 winner, and hence that the high-ability student

[15] If outcomes were binary, and the school with a high outcome wins (ties broken at random) then

$$s=1+\frac{q(G,H)-q(B,H)+[q(G,L)-q(B,L)]}{2}$$

and

$$d=\frac{q(G,H)-q(G,L)+[q(B,H)-q(B,L)]}{2}$$

where $q(i,j)$ denotes the likelihood of a high outcome with school quality $i\in\{G,B\}$ and student ability $j\in\{H,L\}$.

in generation 1 attended the period-0 winning school with probability $\frac{1}{2}+p$. These conjectures imply that the high-ability student in generation 1 is more likely to attend the good school than the bad school, since

$$\begin{aligned} Pr\left(H_{t=1}^{G}\right) &= Pr\left(\overline{G}B_{t=0}\right)\left(\frac{1}{2}+p\right)+\left(1-Pr\left(\overline{G}B_{t=0}\right)\right)\left(\frac{1}{2}-p\right), \\ &= \frac{s}{2}\left(\frac{1}{2}+p\right)+\left(1-\frac{s}{2}\right)\left(\frac{1}{2}-p\right) \\ &= \frac{1}{2}+p\left(s-1\right)\geq\frac{1}{2} \end{aligned}$$

Given this conjecture and given their symmetric prior beliefs, generation 2 updates their beliefs that the period-1 winner is the good school as follows:

$$\begin{aligned} Pr\left(\overline{G}B_{t=1}\right) &= \left[Pr\left(H_{t=1}^{G}\right)\right]\frac{1}{2}\left(s+d\right)+\left[1-Pr\left(H_{t=1}^{G}\right)\right]\frac{1}{2}\left(s-d\right) \\ &= \frac{s}{2}+pd\left(s-1\right)\geq\frac{s}{2} \end{aligned}$$

I can now make the key comparison between posterior beliefs held by generation 2 and generation 1. Generation 2 is better informed than generation 1, despite the fact that each generation only has access to the most recent ranking. The reason is that, relative to the ranking observed by generation 1, the ranking observed by generation 2 is more likely to have been generated in a situation in which the high-ability student attended the good school. That is, the ranking observed by generation 2 is more likely to have been generated in a situation in which the performance bias due to student ability worked in favour of the better school. This is the case because generation 1 had better information about school quality than generation 0. The same logic applies to future generations, and therefore the inference about quality will improve over time.

The fact that generation 1 is better informed than generation 2 can be further illustrated by comparing extreme cases. If an oversubscribed school were to select among students at random ($p=0$) then the high-ability student in generation 1 would be just as likely to attend the better school as the high-ability student in generation 1. Consequently, generation 2 would be as well informed as generation 1. Similarly, if a school's performance were independent of student ability ($d=0$), generation 2 would be as well informed as generation 1. In the case in which performance depends only on student ability ($d=1$ and $s=1$) rankings would be uninformative for either generation.

In the remainder of the paper, I will focus on a steady state in which the share of high-ability students attending the most recent winner is constant across generations. To construct families' posterior beliefs in such an equilibrium, it is not necessary to reason through each generations' choices as in the example above. Instead, it is possible to directly derive the stationary probability that a better school wins, and construct families' posterior beliefs about the relative quality of the winning school based on this stationary distribution.

1.4 Optimal Application Strategy

This section will derive how families in generation t optimally apply to schools, taking as given their posterior beliefs about the relative quality of the period-$t-1$ winning school. How these posterior beliefs are constructed will be the focus of the next section.

It is a dominant strategy for families to apply to the school which they expect to yield a higher payoff if their child were admitted. Each family saves transport costs if their child attends their local rather than their non-local school. Therefore, if families do not expect schools to differ in quality then each family will apply to their local school. By contrast, if families expect schools to differ in quality, then some families find it optimal to apply to the school of higher expected quality even if it is their non-local school. Families in generation t update their beliefs about relative school quality based on the ranking generated in generation $t-1$. Applying to the higher-ranked school, W_{t-1}, yields a benefit of size V if the higher-ranked school is better, and a loss of size V if the higher-ranked school is worse. Therefore, a family with transport costs c who considers W_{t-1} to be their non-local school still applies to W_{t-1} if and only if

$$V\left[\mu\left(\overline{G}B|W_{t-1}\right)-\mu\left(\overline{B}G|W_{t-1}\right)\right] \quad > \quad c$$

where $\mu\left(\overline{G}B|W_{t-1}\right)$ denotes the posterior belief that the higher-ranked school is of better quality, and $\mu\left(\overline{B}G|W_{t-1}\right)$ denotes the posterior belief that the higher-ranked school is of worse quality. Hence, there exists a cut-off level for transport costs, denoted by C_t, such that families whose transport costs lie above the cut-off will find it optimal to apply to their local school, while those whose transport costs lie below

the cut-off will find it optimal to apply to the higher-ranked school, where

$$V\left[\mu\left(\overline{G}B|W_{t-1}\right)-\mu\left(\overline{B}G|W_{t-1}\right)\right]\equiv C_t$$

It is helpful to introduce informativeness as a shortcut for how posterior beliefs enter application choices, and to introduce mobility as a measure of how responsive families' application choices are to rankings.

Definition: ***Period t-informativeness,*** $I_t \in [0,1]$*, is defined as*

$$I_t \equiv \mu\left(\overline{G}B|W_{t-1}\right)-\mu\left(\overline{B}G|W_{t-1}\right)$$

where $\mu\left(\overline{G}B|W_{t-1}\right)$ $\left(\mu\left(\overline{B}G|W_{t-1}\right)\right)$ *denotes generation t's posterior beliefs that the period-t-1 winning school is strictly better (worse) than the other school.*

Definition: ***Mobility of generation*** t*,* $m_t \in [0,1]$*, is defined as the share of families who apply to the period-t-1 winning school,* W_{t-1}*, in the population of families in generation t who consider* W_{t-1} *to be their non-local school.*

Then families' optimal cut-off level translates into a level of mobility (illustrated in Figure 1.2),

$$m_t = F\left(C_t\right) = F\left(V \cdot I_t\right). \tag{1.1}$$

Mobility measures what share of families apply to the winning school rather than their (default) local school which they would have chosen in the absence of information about quality. If families believed schools to be of equal quality, $I_t = 0$, all apply to their local school, $m_t = 0$. If families believed that the period-t-1 winning school is certainly better, $I_t = 1$, then all families whose transport costs fall below the gain from attending a better school apply to the period-$t-1$ winner, $m_t = F\left(V\right)$.

1.5 Equilibrium Steady State

This section will solve for an equilibrium steady state. To understand the equilibrium concept, consider a steady state in which the level of mobility is unchanged over time. I will show that, in steady state,

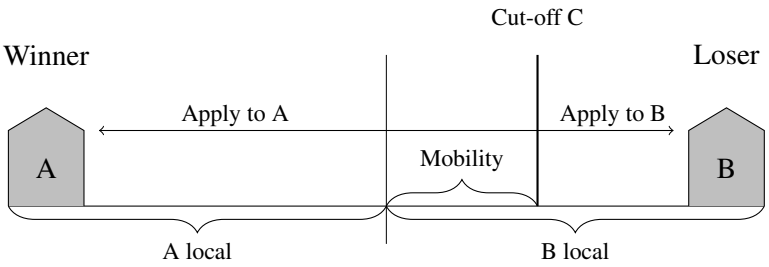

Figure 1.2: Families are distributed along a line between school A and school B. Those with high transport costs live closer to their respective local school. Consider the situation in which school A is the most recent winner. School A receives applications from all families for whom school A is local, and from the share of families for whom school B is local and whose transport costs lie below the cut-off. Mobility is equal to the share of families applying to their non-local winning school.

what families infer about the relative quality of the most recent winner is also unchanged over time, i.e. informativeness is constant. Therefore, if families conjecture to be in a steady state characterised by a constant level of mobility $\widehat{m}$, their level of informativeness is $I(\widehat{m})$, and their optimal application strategies give rise to a level of mobility m:

$$m(\widehat{m}) = F(V \cdot I(\widehat{m}))$$

An *equilibrium steady state* is a steady state such that families behave optimally given their posterior beliefs and beliefs are consistent with the application strategies played by families in earlier generations. For this to be the case, if families were in a steady state characterised by a given level of mobility, their optimal application strategies would need to give rise to this same level of mobility, i.e.

$$m^* = F(V \cdot I(m^*)) \tag{1.2}$$

The remainder of this section will derive how informativeness in steady state depends on the level of conjectured mobility which characterises the steady state, i.e. it will derive $I(\widehat{m})$. I will show that informativeness is a weakly increasing function of conjectured mobility. Since mobility is weakly increasing in conjectured mobility, and since mobility is bounded to lie between 0 and 1, an equilibrium steady-state level of mobility (and informativeness) always exists, but will not necessarily be unique.

Suppose the level of mobility of each generation is equal to $\widehat{m}$ and school qualities are exogenously determined, i.e $\gamma = 0$. For any given school ranking observed by generation t, the level of mobility determines the student allocation in generation t and hence determines how likely it is that generation

$t+1$ observes the same ranking as generation t or the reverse ranking. Therefore, how the distribution of rankings evolves over time can be described by a time-homogeneous Markov process, and its transition probabilities depend on the level of mobility. Eventually, this process will converge to its steady state in which the probability of a given ranking is stationary. Each level of mobility $\widehat{m}$ maps to a unique stationary distribution of rankings. Therefore, if families conjecture to be in a steady state defined by a constant level of mobility $\widehat{m}$, their prior beliefs over school rankings are given by the stationary distribution over rankings when mobility is $\widehat{m}$. Then they update their beliefs about what the current ranking conveys about schools' relative quality using Bayes' rule. Similarly, when school qualities are endogenous, how the distributions of rankings and school qualities evolve can also be characterised by a time-homogeneous Markov process: given the ranking observed by generation t and given the schools' qualities, the level of mobility determines which school is likely to be replaced and how likely it is that generation $t+1$ observes the same ranking as generation t or the reverse ranking. The derivation of the stationary distribution of rankings and the analysis of how families update their beliefs can be found in the Appendix. In each generation, the level of informativeness in steady state as a function of the level of mobility $\widehat{m}$ is given by:

$$I(\widehat{m}) = \frac{2(1-\theta)\theta(s-1)}{(2\theta-1)(s-1)1_{\gamma>0}(\gamma)+1-p\cdot d\cdot\widehat{m}} \tag{1.3}$$

where $1_{\gamma>0}(\gamma)$ denotes an indicator function which is 1 if $\gamma>0$ and 0 otherwise.

Informativeness increases in the level of conjectured mobility as long as $p>0$. This is because, when a school performs better than the other, families not only infer that it is likely to be of higher quality, but also that it is likely to have had a high proportion of high-ability students. This latter factor is informative about what earlier generations thought about the relative quality of this school, since it is the oversubscribed school which enrols a larger share of high-ability students, provided $p>0$. The higher the level of mobility, the higher is the likelihood that the current winner was also the winner in the previous period and hence the more informative is the ranking over time.[16]

For a more technical explanation, imagine school qualities were fixed over time, and one school is good and the other is bad. Consider the transition probabilities between the state in which the good school wins and the state in which the bad school wins (illustrated in Figure 1.3). If conjectured mobility increases, the chances that the period-t-1 winner is also the period-t winner increase for any t. This implies that not only a truthful ranking (good school wins) but also a misleading ranking (bad school

[16]This corresponds to the dynamics illustrated in the example in Section 1.3.

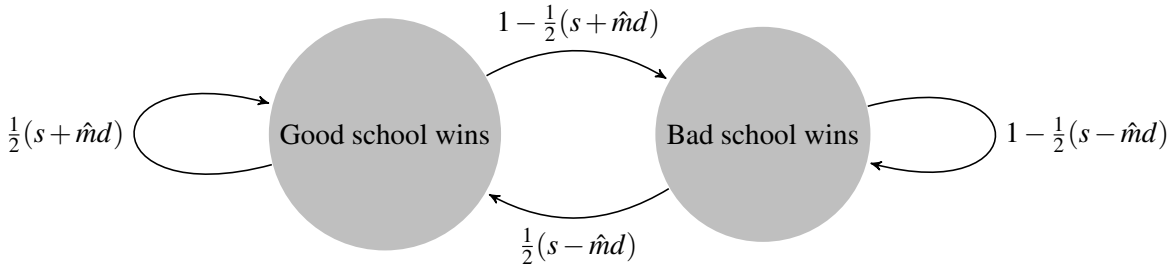

Figure 1.3: Transition probabilities between ranking realizations given conjectured mobility $\widehat{m}$, for the case in which there is one good school and one bad school and school qualities are fixed.

wins) is more likely to be repeated. However, at any given level of mobility, the chances that the period-t-1 winner is also the period-t winner are higher for the good school than the bad school. This is due to the good school's performance advantage $s > 1$. Therefore, in steady state, a higher level of conjectured mobility increases the average number of successive wins by the good school more than it increases the average number of successive wins by the bad school.

Note that if past generations are conjectured to apply to their local school, $\widehat{m} = 0$, or if an oversubscribed school does not admit more high-ability students than the other school, $p = 0$, then student intakes would be the same at each school, irrespective of which school is the most recent winner. The current generation would infer that the winning school is relatively more likely to be the better school, and how much more likely depends on the size of the performance advantage due to superior quality, s. Note also that I assume that the probability of winning is never solely determined by the allocation of students, and therefore the state in which the bad school wins is never an absorbing state. The set-up analysed can be thought of as a case in which schools are sufficiently similar.[17]

Further, note that the level of γ does not affect the steady-state level of informativeness beyond whether it is zero or positive. In a steady state defined by a constant level of mobility $\widehat{m}$, the stationary distribution over pairs of school qualities and their performance-based ranking is independent of the level of γ for $\gamma > 0$. This is because the relative frequency with which a good rather than a bad school is replaced is unaffected by γ, and hence the relative frequency with which a good school ranks higher than a bad school is also unaffected.

Using (1.1) and (1.3) to solve (1.2), the equilibrium steady-state level of mobility and informativeness satisfy the following conditions (and this fixed point of the level of mobility is illustrated in Figure 1.4):

[17] For a discussion about how results are affected when a state could be absorbing see Section 1.9.5

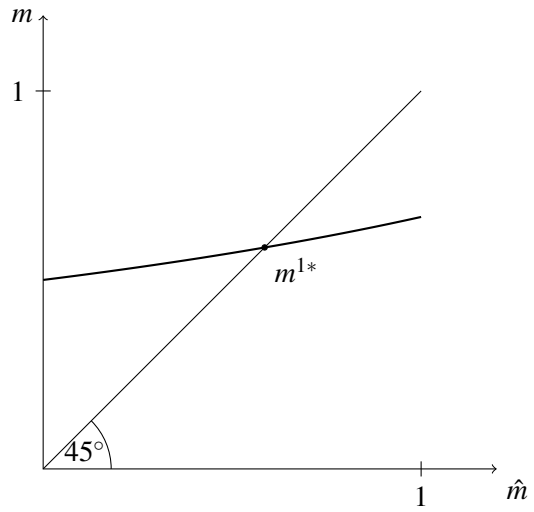

Figure 1.4: The graph shows optimal mobility m as a function of conjectured mobility $\widehat{m}$. The intersection with the 45-degree line shows the steady-state equilibrium level of mobility m^{1*}. This graph is drawn assuming F is a Uniform distribution on $[0,1]$, $V=1$, $\theta=\frac{1}{2}$, $s=\frac{3}{2}$, $p=d=\frac{1}{2}$, and holds for $\gamma\geq 0$.

Proposition 1 [Equilibrium]: *An equilibrium steady-state level of mobility is characterised by*

$$m^{*}=F\left(V\cdot\frac{2\left(1-\theta\right)\theta\left(s-1\right)}{\left(2\theta-1\right)\left(s-1\right)1_{\gamma>0}\left(\gamma\right)+1-p\cdot d\cdot m^{*}}\right) \tag{1.4}$$

and the corresponding equilibrium steady-state level of informativeness is given by

$$I\left(m^{*}\right)=\frac{2\left(1-\theta\right)\theta\left(s-1\right)}{\left(2\theta-1\right)\left(s-1\right)1_{\gamma>0}\left(\gamma\right)+1-p\cdot d\cdot m^{*}}. \tag{1.5}$$

where $1_{\gamma>0}\left(\gamma\right)$ *takes value 1 if* $\gamma>0$ *and 0 otherwise.*

Such an equilibrium steady-state level of mobility (and informativeness) always exists.

1.5.1 Convergence to Equilibrium Steady State

In what follows, I will focus on the smallest equilibrium steady-state level of mobility and informativeness. This is because these levels of mobility and informativeness are approached in the limit, assuming each generation of families responds optimally to the mobility levels of all previous generations.

Proposition 2 [Convergence] *For $\gamma = 0$, the dynamic learning process $\{I_t\}_{:t=0,1,...}$ is characterised by the informativeness of generation 0 who are the first generation to apply to schools,*

$$I_0 = 0$$

and by the recurrence equation for informativeness which arises if each generation responds optimally to the behaviour of all previous generations,

$$I_t = 2(1-\theta)\theta(s-1) + F(V \cdot I_{t-1})\, p \cdot d \cdot I_{t-1}.$$

Its limit as $t \to \infty$, denoted by I^{1*}, is the smallest equilibrium level of informativeness. As $\{I_t\}_{:t=0,1,...}$ converges so does the sequence of mobility levels $\{m_t\}_{:t=0,1,...}$. The sequence $\{m_t\}_{:t=0,1,...}$ converges to the smallest equilibrium level of mobility, denoted by m^{1*}, where

$$m^{1*} = F\left(V \cdot I^{1*}\right).$$

If school qualities are exogenously determined, i.e. $\gamma = 0$, starting from any level of informativeness, the sequence of levels of informativeness that arises from families' optimal application strategies is increasing and will converge to an equilibrium level of informativeness. Starting from a symmetric prior across schools, i.e. a level of informativeness equal to zero, such a sequence will converge to the smallest equilibrium level of informativeness. As equilibrium levels of informativeness and mobility are jointly ordered, this also implies that the sequence of mobility levels will converge to the smallest equilibrium level of mobility.[18] If school quality is endogenous, i.e. $\gamma > 0$, the learning process is described by a more complex system of recurrence equations. Simulations support the conjecture that the sequences of levels informativeness and mobility still converge to their smallest equilibrium steady-state levels (for detail see Appendix).

[18] Note that even if families do not respond optimally to past generations' mobility, but simply apply to the winning school with a fixed probability, then the sequence of levels of informativeness converges. However, it may not converge to an equilibrium steady-state level.

1.6 Comparative statics

This section will analyse how families' inferences and optimal application strategies in equilibrium steady state vary across different environments. The following analysis will focus on the equilibrium steady state associated with the smallest level of mobility (m^{1*}) and the corresponding level of informativeness (I^{1*}), henceforth equilibrium levels of informativeness and mobility (see Section 1.5.1).

Proposition 3 [Comparative Statics]

- *An increase in the proportion p of non-local high-ability applicants admitted by an oversubscribed school, or*
- *a negative shift in FOSD of the distribution of transport costs, $F(\cdot)$,*

weakly increases the equilibrium levels of mobility and informativeness if $\gamma \geq 0$, and weakly increases the average fraction of good schools, denoted by θ, if $\gamma > 0$.

Exogenous School Qualities **-** First, I will assume school qualities are exogenously determined and fixed over time, i.e. $\gamma = 0$.[19]

Both of these changes cause optimal mobility to weakly increase at any given level of conjectured mobility. Hence, the smallest fixed point weakly increases (see Figure 1.5).[20] Optimal mobility increases because more high-ability students are admitted at the winning school for any given cut-off level of transport costs - either because i) the oversubscribed winning school admits a larger proportion p of high-ability applicants from any given applicant pool, or ii) because a negative shift in the distribution for transport costs causes an increase in the applicant pool and therefore an increase in the proportion of high-ability students admitted provided $p > 0$.[21] If more high-ability students are admitted by the winning school, families are more confident that a higher rank is an indication that past generations believed this school to be of higher quality and hence more families are willing to apply to the higher-ranked school.[22]

[19] If instead of holding school quality fixed, I assumed that an exogenously chosen school was replaced by a new school with a probability strictly less than 1, the analysed effect would still exist but be weaker.

[20] The equilibrium levels of informativeness and mobility also weakly increase with a rise in the performance advantage due to superior quality, s, or with a rise in the performance advantage due to stronger students, d. Just as a negative shift in FOSD of the distribution $F(\cdot)$, an increase in s shifts up optimal mobility at any given level of conjectured mobility. And just as an increase in p, an increase in d shifts up optimal mobility at any positive level of conjectured mobility.

[21] This effect also arises if families' payoff V from attending a good school increases.

[22] This intuition is analogous to the one explained in Section 1.5 for why mobility increases in the level of conjectured mobility of past generations.

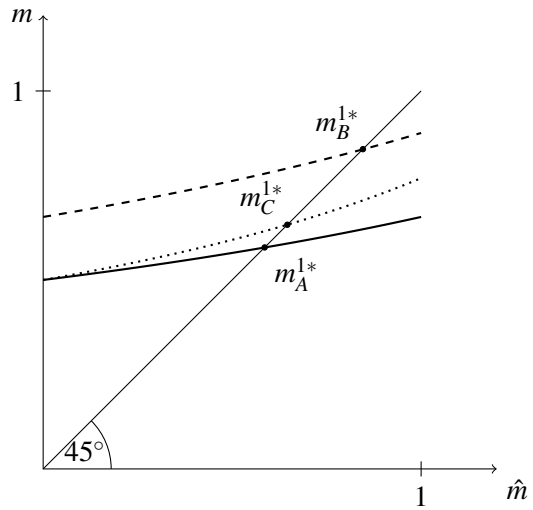

Figure 1.5: Consider the equilibrium level of mobility m_A^{1*}. A negative shift in FOSD of the distribution $F(\cdot)$ shifts up optimal mobility at any level of conjectured mobility, as illustrated by the dashed line. The equilibrium level of mobility increases to m_B^{1*}. An increase in p shifts up optimal mobility at any positive level of conjectured mobility, as illustrated by the dotted line. The equilibrium level of mobility increases to m_C^{1*}.
The solid line is drawn for F being a Uniform distribution on $[0,\overline{c}]$, $\overline{c}=1$, $V=2$, $\theta=\frac{1}{2}$, $s=\frac{3}{2}$, $p=d=\frac{1}{2}$, and holds for $\gamma\geq 0$. The dashed line is drawn for $\frac{V(s-1)}{\overline{c}}=\frac{8}{3}$, and the dotted line for $p\cdot d=\frac{7}{20}$, all else equal.

These results refute the claim that the use of admission lotteries would make it easier to identify good schools from performance in a setting in which families have limited knowledge of past performance. Instead, performance-based rankings are more likely to reflect schools' underlying relative quality if oversubscribed schools allocate a larger share of places based on measures of student ability. This highlights the novel insight that assigning priorities in admission codes based on certain student characteristics such as location of residence can have negative consequences for the aggregation of information about school quality over time.

In addition, the results demonstrate that families are more likely to identify higher-quality schools over time if it is less costly to choose a school based on its quality. For example, the introduction of school choice has lowered families' costs associated with sending their child to a non-local school. Families are no longer required to move to this school's attendance area; instead they can simply commute. This should facilitate learning about schools' relative quality. The costs c of choosing the non-local school represent the costs of choosing one school over another if both were of the same quality. Therefore, the result more generally implies that, if schools are similar in other respects, e.g. if schools specialise in similar areas, better schools are more easily identified over time.

The costs c can also capture the time and effort incurred by families to research how schools compare in the latest performance-based ranking. Unlike the transport costs, this would be incurred independent of families' application choice. Families then trade off applying to their local school without doing any research against incurring the costs and applying to the better-performing school. The results suggest that the more easily accessible performance-based rankings are, the more informative these rankings will become over time.[23].

Endogenous School Qualities - I will introduce the feature that an undersubscribed school is replaced by a new school with probability ($\gamma > 0$) to incorporate that a school which cannot attract applications is under pressure to change or close down. The replacement of the school does not necessarily represent a school closing down, but could also capture that a school changes its administration or hires new teachers. The new school's quality is drawn from the same pool as the original schools' qualities. Therefore, an undersubscribed school does not necessarily improve when it is replaced. Yet, the composition of school qualities is affected by demand, because families' demand determines whether the better or the worse of the two schools will be replaced. If families are more likely to demand the better school, high-quality schools operate for longer than low-quality schools and the average quality increases. Because I focus on the equilibrium steady state, the probability γ with which an undersubscribed school is replaced does not affect outcomes, provided $\gamma > 0$. Intuitively, the quality of schools in equilibrium steady state depends on the relative frequency with which a good rather than a bad school is replaced, and this is unaffected by γ for $\gamma > 0$.[24]

Deriving the comparative statics is not as straightforward as when school qualities are exogenously determined, because I need to take into account how changes in school qualities in turn affect families' inferences about schools' relative quality and hence families' application choices. If the recent winner admits more high-ability applicants at any level of conjectured mobility, then the equilibrium levels of mobility and informativeness weakly increase, as in the case of exogenously determined school qualities. However, higher informativeness now also triggers a rise in the average fraction of good schools. If the fraction of good schools in the market is smaller than $\frac{1}{2}$, the variation in school qualities increases.[25] As families conjecture that schools are more likely to be of different quality, they are more willing to travel

[23]That access to performance information can increase applications to well-performing schools has been shown by Hastings and Weinstein [2007]

[24]In particular, γ only affects the rate of convergence to steady state but not the limit the process converges to.

[25]This is the case for θ sufficiently small. Note that since school quality is binary changes in mean and variance are linked.

to the better-performing school. Hence, mobility increases by more than in the case of exogenously determined qualities at any given level of conjectured mobility. By contrast, if the fraction of good schools in the market is larger than $\frac{1}{2}$, variation in school qualities decreases. Hence, mobility increases by less than in the case of exogenously determined qualities.

These results closely match empirical findings in health care choice by Chandra et al. [2016]. They find that higher-quality hospitals have higher market shares, and that this is particularly pronounced for patients who do not access the hospital via the emergency room and hence have greater scope for choice. In my model, higher scope for choice is captured by lower transport costs and Proposition 3 shows that the average quality of schools is higher when transport costs are lower.

A natural comparative static exercise is to ask how the parameter γ affects the outcomes of interest when it increases from zero to a positive value. One could interpret this change in γ as the introduction of educational reforms which put pressure on undersubscribed schools to improve either by withdrawing funding or by forcing these schools to take measures for improvement.

Proposition 4 [Exogenous vs. Endogenous Quality]

If an undersubscribed school is replaced by a new school, i.e. $\gamma > 0$ rather than $\gamma = 0$, then this

- *weakly increases the equilibrium levels of informativeness and mobility if and only if $\theta \leq \frac{1}{2}$,*
- *weakly increases the average fraction of good schools.*

If γ changes from $\gamma = 0$ to $\gamma > 0$, the average fraction of good schools increases at any given level of conjectured mobility. The effect on equilibrium levels of mobility and informativeness depends on the quality of incoming schools. If a new school is likely to be good, i.e. $\theta \geq \frac{1}{2}$, then applying to the most recent winner does not improve the likelihood of attending a better school by as much as it would if school qualities were fixed. Therefore, fewer families apply to the most recent winner than prior to the change in γ, at any given level of conjectured mobility. By contrast, if this new school is likely to be bad, i.e. $\theta \leq \frac{1}{2}$, then more families will apply to the most recent winner than prior to the change in γ. Hence, the equilibrium levels of mobility and informativeness increase if and only if $\theta \geq \frac{1}{2}$. Whenever the equilibrium level of informativeness increases, the fraction of good school rises more strongly than when

informativeness decreases. This is because at a higher level of informativeness it is even more likely that the undersubscribed school is the worse of the two schools. These findings suggest that putting pressure on undersubscribed schools improves the average quality of schools in the long run, and this improvement is stronger (weaker) if the variance in schools' qualities increases (decreases).

1.7 Welfare

This section studies how changes in schools' admission policies and in the distribution of transport costs affect the allocation of students in steady state. In addition, I will also discuss how these changes impact on welfare.

Proposition 5 [Access to Good Schools - Selection Criteria]

- *An increase in the proportion p of non-local high-ability applicants admitted by an oversubscribed school, or*
- *a negative shift in FOSD of the distribution of transport costs, $F(\cdot)$,*

weakly increases the share of high-ability students who attend a good school and weakly decreases the share of low-ability students who attend a good school for $\gamma \geq 0$.

Exogenous School Qualities - If school qualities are exogenously determined then any change which increases the share of high-ability students at a good school must also decrease the share of low-ability students at a good school.

If oversubscribed schools take on a larger share of high-ability students from a given applicant pool, then the share of high-ability students attending a good school will increase. This effect is stronger in a dynamic setting than it would be in a static setting, because over time the change in the allocation of students will also increase the likelihood that the oversubscribed school is the better of the two schools. From a utilitarian social planner's point of view, welfare decreases, because the overall number of good schools remains unchanged while the level of transport costs incurred increases. On average, high-ability students are better off and low-ability students are worse off.

A policy which lowers transport costs also strictly increases the share of high-ability students attending a good school, provided $p > 0$. But it does not necessarily increase the overall expenditure on

transport costs: even if more families travel to their non-local school it could be that those who travel incur a lower average cost.

Endogenous School Qualities - If an undersubscribed school is replaced by a new school with some probability, then changes in schools' admission policies and changes in families' transport costs have the potential to affect the average quality of schools and improve access to good schools for both high- and low-ability students.

Students' access to good schools can be thought of as being subject to two separate effects: one is the effect on how many of the schools are good on average (quality effect), and the other is the effect on how many high-ability students are allocated to the relatively better of the two schools conditional on schools being of different quality (sorting effect). While an improvement in school quality benefits both types of students, a stronger sorting effect benefits only high-ability students, and is detrimental for low-ability students.

A higher level of mobility increases both the fraction of good schools and the degree of sorting between student ability and school quality. Therefore, the average high-ability student is better off. However, the share of low-ability students at good schools decreases, despite the improvement in the average quality of schools, because the sorting effect outweighs the quality effect. Consequently, the average low-ability student is worse off if an oversubscribed school accepts a larger share of high-ability applicants. He may be better off after a negative shift in FOSD of the distribution of transport costs if this reduces his expenditure sufficiently to outweigh the loss from attending a good school less frequently. From a utilitarian social planner's point of view, the impact of more selection on welfare depends on the extent to which transport costs increase relative to the increase in the average fraction of good schools, and hence depends on the functional form of F. The proposition shows that it is not necessarily the case that policy changes which increase the informativeness of rankings also improve access to good schools for all students. In addition, it points out that there may exist a trade-off between equity and efficiency. On the one hand, enabling better-performing schools to admit more able students has the advantage that rankings are more likely to reflect schools' relative quality over time, which causes worse schools to face more pressure to improve. On the other hand, high-ability students benefit disproportionately from the resulting improvement in school quality, because they are more likely to be admitted to a relatively better school.

Finally, a change in γ from zero to a positive value also affects students' access to good schools.

Proposition 6 [Access to Good Schools - Exogenous vs Endogenous Quality]

If an undersubscribed school is replaced by a new school, i.e. $\gamma > 0$ *rather than* $\gamma = 0$*, then this*

- *weakly increases the share of high-ability students who attend a good school if*

$$\theta < max\left\{\frac{1}{2}, \frac{1-p(2+d-s)}{2p(s-1)}\right\},$$

and

- *weakly increases the share of low-ability students who attend a good school.*

Introducing replacement of schools raises the average quality of schools at each level of mobility. If the average quality of incoming schools is lower than the average quality in the market, then both sorting and the average quality of schools are higher in the new equilibrium steady state. The increase in the average quality of schools is strong enough such that the share of low-ability students at good schools increases. However, since the average low-ability student incurs higher transport costs than before, the effects on the average low-ability student's payoff are ambiguous. By contrast, the average high-ability student is unambiguously better off.

If the average quality of incoming schools is higher than the average quality in the market, then the average quality of schools is higher and sorting is lower in the new equilibrium steady state. Low-ability students benefit both from a higher average quality of schools and from decreased sorting. Since their expenditure on transport costs also decreases, they are unambiguously better off. For high-ability students, the benefit of higher average quality is not always sufficiently strong to outweigh the negative consequences of decreased sorting such that the average high-ability student may attend a worse school. This does not necessarily mean the average high-ability student is worse off because he also spends less on transport. From a utilitarian planner's point of view, welfare increases unambiguously. However, the cost associated with replacing schools is not incorporated in this model and may outweigh the benefit discussed here.

1.8 Externalities

This section will highlight the externalities arising between families within the same generation and across generations. I will analyse how a utilitarian social planner would allocate students to schools if the social planner was subject to the same informational constraints as families.

In particular, assume a social planner commits to a rule which assigns families to schools in each generation with the objective to maximise utilitarian welfare in steady state. In each generation, he can condition this assignment on families' transport costs, on the ability of their child, and on the signal about school qualities available to the current generation.

First, consider a static setting with fixed school qualities to analyse externalities within a given generation. Each high-ability student who attends his non-local school takes up a place that would have otherwise been available for a low-ability student who lives closer to this school. A utilitarian social planner would consider this socially wasteful and would prefer each student to attend their respective local school.

Second, consider a dynamic setting to analyse externalities across generations which arise because families infer schools' qualities from the most recent ranking. Each family weighs up the expected increase in payoff from attending a better-performing school against their level of transport costs. However, what they do not take into account is that if more families applied to the better-performing school, then informativeness would increase and good schools would stay in the market relatively longer than bad schools. Their application decision therefore has an influence on the average quality of schools available for future generations of families. In this setting, a utilitarian social planner has a reason to incur transport costs in order to allocate some non-local high-ability students to the better-performing school, because he benefits from the rise in the average quality of schools in steady state.

It is worth noting that such an externality would persist even if families did not make application choices strategically. Suppose a certain share of families applies to the winning school, whereas the remainder applies to their local school, irrespective of the level of informativeness. If a larger share of families in past generations applied to the winning school, then families in the current generation are better off. This is because if a larger share of families in past generations applied to the winning school, then a better school was more likely to be oversubscribed, and hence the schools available for the current generation have a higher average quality.

To illustrate both these externalities, I solve for the share ν of high-ability students a social planner

would assign to the better-performing school to maximise utilitarian welfare in steady state. I will assume that conditional on ability he will assign those students with the lowest transport costs to attend their non-local schools. His optimisation problem is described as follows:

$$\max_{v\in[0,1]} W =$$
$$\max_{v\in[0,1]} 2\left[V\frac{\theta\left(\theta\left(s-1\right)+1-dv\right)}{\left(2\theta-1\right)\left(s-1\right)+1-dv}-\int_0^{F^{-1}(v)} cf\left(c\right)dc\right]$$

where F^{-1} denotes the inverse of F.

Whether the externality within a given generation or across generations is stronger, and hence whether the social planner allocates a larger or smaller share of high-ability students to the better-performing school relative to the share in equilibrium depends on the distribution of transport costs. If the distribution for transport costs is Uniform on $[0,\overline{c}]$ then there exists a critical value of $\overline{c}$ such that for $\overline{c}$ below the critical value, the social planner would assign all high-ability students to the better-performing school. At the critical value, the share of high-ability students assigned to the better-performing school is reduced discontinuously, and for $\overline{c}$ above the critical value, the share of high-ability students is reduced continuously with $\overline{c}$.

1.8.1 Complements or Substitutes

From a social planner's point of view, there may be complementarities or substitutes between student ability and school quality. If there are complementarities (substitutes) between student ability and school quality, then even in a static setting a social planner may be willing to incur transport costs to increase the share of high-ability (low-ability) students attending the better-performing school. This is because the additional expenditure on transport costs could be outweighed by the gain derived from reassigning students such that those who benefit more from a better school attend the school of higher expected quality.

In a dynamic setting, the presence of complementarities would lead a social planner to assign an even larger share of high-ability students to the better-performing school than in the static setting. This is because as a larger share of high-ability students is assigned to the better-performing school, the likelihood that this school is good rises, which further increases incentives to assign high-ability students to this school. However, in the case of substitutes, a trade-off arises. Welfare of the current generation is

higher if more low-ability students than high-ability students attend the better-performing school. However, assigning more low-ability students to the better-performing school lowers the overall quality of schools in steady state. Even if there was no replacement of schools, i.e. $\gamma = 0$, then allocating more low-ability students to the better-performing school in the current generation reduces the chances that the better school can be identified in future generations, and hence it will reduce the chances that low-ability students will attend the better of the two schools in future generations.

1.9 Extensions and Limitations

1.9.1 Longer Window of Performance-based Rankings

To obtain the results of this paper, it is not crucial that families observe only the most recent ranking of schools. If families observed a longer window of rankings then they would have access to some of the information based on which previous generations applied to schools. However, they still would not have access to all the information available to previous generations. Therefore, if families in each generation applied to the school which they expect to be of higher quality then some additional information about the relative quality of schools would be conveyed from one generation to the next. Hence, informativeness would still increase in the conjectured levels of mobility of past generations.

In the Appendix, I analytically solve for the case in which families observe the two most recent rankings, assuming school qualities are exogenously determined. Families' optimal mobility level is allowed to differ depending on whether or not they observe that the same school won in both rankings. I show that in steady state it is uninformative for families to observe a different school win in each ranking, independent of the level of conjectured mobility. By contrast, observing the same schools win in both rankings is informative, and weakly more so the higher the level of conjectured mobility of past generations.[26] In Section 1.11.7 in the Appendix, I also report the results of simulations for the case in which school qualities are endogenous, which also show that informativeness is increasing in mobility.

[26] Since families can condition their application choices on whether or not they observe the same school win in both rankings, application strategies in a given generation are characterised by two levels of mobility. Observing different schools win is more informative about schools' relative quality, the higher the level of conjectured mobility characterising choices when different schools were observed to have won, and it is independent of the level of conjectured mobility characterising choices when the same school was observed to have won.

1.9.2 No Ability-based Selection by Schools

In many settings, schools cannot select among applicants based on their ability, or are forced to prioritise local applicants over non-local applicants. In the benchmark model, an increase in mobility strictly increases informativeness only if the better-performing (and hence oversubscribed) school admits some non-local high-ability applicants, i.e. $p > 0$. Yet the basic insight of the model still applies in situations in which schools cannot select based on ability, provided i) the school of higher expected quality admits a larger share of high-ability students, and ii) provided this share of high-ability students weakly increases the stronger the belief that this school is of higher quality.

Assume that an oversubscribed school has to use a lottery to assign places to applicants. If high-ability applicants are more mobile than low-ability applicants, then the better-performing school may still admit a larger share of high-ability applicants than the other school. In fact, there is empirical evidence that students of higher socio-economic status are more likely to seek out better-performing schools.[27] This could be because high-ability applicants tend to value attending a good school more, or tend to be better informed about schools' most recent performance, or tend to have lower transport costs. This could be incorporated into the model by assuming a positive correlation between the payoff V from attending a better school and ability H. In this setting, families' application strategies would not only depend on their transport costs but also on their child's ability. The main insights of the benchmark model still apply if the distribution of transport costs satisfies the following condition: for any given increase in the level of informativeness the increase in the share of applications from high-ability students is larger than the increase in the share of applications from low-ability students.

In addition, it is then possible for a policy measure to increase willingness to travel to better-performing schools in the overall population of students and yet lower informativeness. This is the case if a measure disproportionately increases the mobility of low-ability students relative to high-ability students.[28] While such a policy would create more equal access to good schools across student types, it would also make it harder for future generations to identify which school is better, which has detrimental effects on the overall quality of schools in the long run. (See Section 1.11.8 in the Appendix for a detailed analysis).

Even if a school had to admit those students who live nearest to the school, some of the effects captured in this model may still be present. Families would consider the quality of the school when

[27] E.g. see Hastings and Weinstein [2007], Allen et al. [2014]

[28] Such scenarios are studied in Gavazza and Lizzeri [2009], but they do not study the impact on school rankings.

choosing where to live (Tiebout choice). If families with higher-ability children are more likely to move into the proximity of the better-performing school, either because they value school quality more or because they are more likely to be able to afford housing there, then again a better-performing school may admit more high-ability students.

1.9.3 Peer Effects

In the context of schools, families may not only care about the quality of the school but also about the ability of their child's peers. Families' inference problem, which is the focus of this paper, would still persist even in the presence of such peer effects. Schools' past performance is influenced by students who attended these schools in the past, not by the child's peers. If I incorporate concerns for peers into the model and analyse the effects on the equilibrium steady state, the findings of the baseline model change in two important ways. First, the smallest equilibrium levels of mobility and informativeness increase relative to the baseline model which demonstrates that peer effects can reinforce the increase in informativeness and sorting described in the baseline model. Second, unlike in the baseline model, a negative equilibrium level of mobility may exist.[29]

To keep the exposition simple, I will assume that there is one good and one bad school, and their qualities are fixed over time. Families know that one school is better, and have symmetric prior beliefs about which school this is. I will capture peer effects by assuming that families derive a benefit $V > 0$ in the event that the school their child attends wins, rather than in the event that this school is good. The payoff of families in generation t is now given by:

		Performance in period t	
		winner	loser
Proximity	local	V	0
	non-local	$V-c$	$-c$

Consequently, their payoff increases in both the school's quality and the ability of the current intake, and is independent of their child's ability. Parameter d can be interpreted to capture the strength of peer effects. The stronger peer effects, the stronger families' preference for the school with a higher number of high-ability students, all else equal.

[29]Rothstein [2006] examines the relative valuation of peers versus observed school effectiveness in Tiebout choice, and points out that a multiplicity of equilibria arises if families value peers, in some of which families do not sort according to school effectiveness. The negative level of equilibrium mobility which can arise in my model has a similar flavor, but my model also shows that peer effects can strengthen the informativeness of school performance rankings.

Proposition 7 [Peer Effects] *The smallest positive equilibrium levels of mobility and informativeness increase as i) peer effects are stronger (d larger), or as ii) an oversubscribed school selects a larger proportion p of non-local high-ability applicants. Further, in the presence of peer effects, some share of families may apply to the non-local losing school, i.e. a negative equilibrium level of mobility may exist, and it is more likely to exist as i) peer effects are stronger (d larger), or as ii) an oversubscribed school selects a larger proportion p of non-local high-ability applicants.*

The smallest positive equilibrium steady-state levels of mobility and informativeness are higher in the presence of peer effects. This is because if families expect a larger share of high-ability students to apply to the winning school, then this raises the benefit of attending this school for any given level of informativeness. As a larger share of families applies to the winning school, informativeness rises by more, reinforcing the initial effect. Stricter selection on ability or stronger peer effects unambiguously raise the equilibrium level of mobility and informativeness.

The introduction of peer effects can give rise to steady-state equilibria associated with a negative level of mobility, i.e. families incur costs to travel to the *losing* school. Families still assign a higher belief to the winning school being good. Yet if they expect the share of high-ability students at the losing school to be sufficiently high, the expected benefit derived from better peers at a worse school may outweigh the expected benefit derived from a better school with worse peers. These equilibria are more likely to arise if peer effects are stronger (d higher) or if oversubscribed schools admit an even larger share of high-ability students (p higher).

1.9.4 Signal of School Quality

The following two properties of the signal are sufficient for the results of this paper:

Denote by $P^{GB}(h)$ the probability that the good school wins if it takes on a share h of high-ability students, and denote by $P^{BG}(h)$ the probability that the bad school wins if it takes on a share h of high-ability students.

1) The good school has an intrinsic performance advantage for all h, i.e.

$$P^{GB}(h) - P^{BG}(h) > 0. \tag{1.6}$$

2) The increase in the chances that the good school wins when an additional high-ability student is added to the good school's intake comprising a share h of high-ability students is larger than the increase in the chances that the bad school wins when an additional high-ability student is added to the bad school's intake comprising a share h of high-ability students for all h, i.e.

$$\frac{\partial}{\partial h}P^{GB}(h) \geq \frac{\partial}{\partial h}P^{BG}(h). \tag{1.7}$$

The functional form used in this model is chosen such that the extent to which performance is biased due to the allocation of students is symmetric across schools and can be measured by parameter d:

$$\frac{\partial}{\partial h}P^{GB}(h) = \frac{\partial}{\partial h}P^{BG}(h) = d > 0.$$

In addition, the extent to which superior quality gives rise to a performance advantage can be measured by parameter s:

$$P^{GB}(h) - P^{BG}(h) \quad = \quad s - 1 > 0.$$

If I assumed noise entered additively then in order to satisfy (1.6) and (1.7) I would need to assume sufficiently strong complementarities between student ability and school quality in performance (for details see Section 1.11.10 in the Appendix).

1.9.5 Absorbing States

The paper finds that informativeness of rankings is higher if oversubscribed schools admit more high-ability students on average than if student intakes are balanced across schools. Would this finding still hold even if student intakes became very imbalanced, so imbalanced that the school with better students were guaranteed to perform better? In particular, assume the following change to the signal: there exist a lower and an upper critical value for the share of high-ability students at the good school h, denoted by $\underline{h}$ and $\overline{h}$ respectively, such that if $h > \overline{h}$, $Pr\left(\overline{G}B\right) = 1$ and if $h < \overline{h}$, $Pr\left(\overline{G}B\right) = 0$, and otherwise $Pr\left(\overline{G}B\right) = \frac{s}{2} + \left(h - \frac{1}{2}\right)d$.

If oversubscribed schools select on ability, a self-perpetuating cycle could arise in which the same school continues to win, because once it has attracted sufficiently many high-ability students to win it

will then be able to attract these students again and again. The process would enter an absorbing state. It is possible that the bad school is the first to attract the critical mass of high-ability students. Then the level of informativeness in the long run would be lower than in the case of balanced intakes. However, it is more likely that a good school rather than a bad school will be the first to admit sufficiently many high-ability students. Therefore, from an ex-ante perspective, informativeness is still at least as high if oversubscribed schools select on ability as if student intakes are balanced across schools.

The paper has abstracted away from such absorbing states to keep the analysis tractable. The model can be thought of as capturing situations in which there is sufficiently high differentiation between schools such that intakes will not become too unbalanced across schools, or situations in which students' abilities are relatively homogeneous so that the advantage from taking on stronger students is not too large.

1.9.6 Replacement of Unpopular Schools

The model assumes that undersubscribed schools are replaced by new schools. In addition, the quality of the entering school is assumed to be independent of the quality of the exiting school. However, I could instead assume that the quality of the entering school is correlated with the quality of the exiting school. As long as the entering school is not good or bad with certainty, the key result that stricter selection by schools increases informativeness and the average quality of schools would persist.

The assumption that school types are binary, however, is restricting the analysis. Any increase in the mean quality of schools raises the variance if the mean is low, and lowers the variance if the mean is high. For a more general distribution, this link between mean and variance need not exist. Informativeness depends on the variance in school qualities. Therefore, it is not sufficient to look at changes to the mean of school qualities to analyse how informativeness behaves, but changes to the variance would also need to be considered.

1.10 Conclusion

This paper studies families' inferences about school quality from rankings, when families in each generation observe which of two schools performed better with students in the previous cohort, but they do not observe past student allocations or application decisions. I develop a dynamic framework which allows a tractable analysis of families' inferences about school quality and analyse comparative statics in an equi-

librium steady state. I find that a performance-based ranking is more informative about school quality if oversubscribed schools select a larger share of their intake based on ability. I also find that such a ranking is more informative if it is less costly for families to attend a non-local school. My findings contribute to a better understanding of how families choose between schools, and in particular which factors determine how responsive families' demand is to the underlying quality of schools. These findings are important in the light of recent discussions on school choice and on the design of school admission codes.

I view the framework developed as a building block for future research on analysing the link between information and match outcomes. An interesting extension would be to analyse how various matching algorithms affect learning about schools' quality in a set-up with multiple schools. Depending on the mechanism, students or schools may have incentives to act strategically in their application and admission choices. In particular, families may find it optimal to express a preference for a school to which their child is likely to be admitted (safe school) rather than express a preference for the one with the highest perceived quality.[30] These strategic applications would interact with the inferences made by future generations in interesting ways. In addition, the framework is also suitable to explore what could cause persistent differences between schools to arise, and a starting point to explore how schools build and maintain reputations over time when families have limited access to past performance information.

1.11 Appendix

1.11.1 Equilibrium Steady State: Proposition 1

Characterisation - To complete the proof of why the equilibrium is characterised by (1.4) and (1.5), derive

$$I(\widehat{m}) = \frac{2(1-\theta)\theta(s-1)}{(2\theta-1)(s-1)1_{\gamma>0}(\gamma)+1-p\cdot d\cdot\widehat{m}}. \tag{1.3}$$

Families use Bayes' rule to update their beliefs on how likely it is that the school they observe winning is a better school rather than a worse school. They conjecture that the process generating rankings is in a steady state, and the level of conjectured mobility of past generations $\widehat{m}$ determines the stationary

[30] Calsamiglia and Güell [2014] find empirical evidence that, in a system with a Boston mechanism, families base their preferences over schools on how high a priority their child is assigned at each school. In particular, they found that many families switched their first choice school after neighbourhood boundaries, and hence priorities, were reassigned.

distribution defined by the vector of probabilities,

$$\left(\rho^{GG}(\widehat{m}),\rho^{\overline{G}B}(\widehat{m}),\rho^{\overline{B}G}(\widehat{m}),\rho^{BB}(\widehat{m})\right), \tag{1.8}$$

where $\rho^{GG}(\widehat{m})$ denotes the stationary probability that both schools are good, $\rho^{\overline{G}B}(\widehat{m})$ denotes the stationary probability that schools are of different quality and the good school wins, $\rho^{\overline{B}G}(\widehat{m})$ denotes the stationary probability that a schools are of different quality and the bad school wins, and $\rho^{BB}(\widehat{m})$ denotes the stationary probability that both schools are bad. Denote by i the identity of the winning school and j the identity of the losing school, then posterior belief by families in generation t that school i is the better school is given by

$$\mu\left((i,j)=(G,B)\,|i=W_{t-1}\right)=\frac{\rho^{\overline{G}B}(\widehat{m})}{\rho^{GG}(\widehat{m})+\rho^{\overline{G}B}(\widehat{m})+\rho^{\overline{B}G}(\widehat{m})+\rho^{BB}(\widehat{m})}.$$

Due to the symmetry of the set-up, their updated beliefs do not depend on the identity of the school they observe winning. Therefore, I write

$$\mu\left(\overline{G}B|W_{t-1}\right)=\rho^{\overline{G}B}(\widehat{m}). \tag{1.9}$$

Similarly,

$$\mu\left(\overline{B}G|W_{t-1}\right)=\rho^{\overline{B}G}(\widehat{m}). \tag{1.10}$$

By definition of informativeness and (1.9) and (1.10):

$$\begin{aligned} I_t(\widehat{m}) &\equiv \mu\left(\overline{G}B|W_{t-1}\right)-\mu\left(\overline{B}G|W_{t-1}\right) \\ &= \rho^{\overline{G}B}(\widehat{m})-\rho^{\overline{B}G}(\widehat{m}). \end{aligned} \tag{1.11}$$

To derive the stationary distribution described by (1.8), denote by $\overline{kl}_t$ the event in which the period-$t-1$ winner is of quality k and the loser is of quality l where $l,k\in\{G,B\}$. Define the transition matrix T as

$$T = \begin{pmatrix} Pr\left(GG_t|GG_{t-1}\right) & Pr\left(\overline{G}B_t|GG_{t-1}\right) & Pr\left(\overline{B}G_t|GG_{t-1}\right) & Pr\left(BB_t|GG_{t-1}\right) \\ Pr\left(GG_t|\overline{G}B_{t-1}\right) & Pr\left(\overline{G}B_t|\overline{G}B_{t-1}\right) & Pr\left(\overline{B}G_t|\overline{G}B_{t-1}\right) & Pr\left(BB_t|\overline{G}B_{t-1}\right) \\ Pr\left(GG_t|\overline{B}G_{t-1}\right) & Pr\left(\overline{G}B_t|\overline{B}G_{t-1}\right) & Pr\left(\overline{B}G_t|\overline{B}G_{t-1}\right) & Pr\left(BB_t|\overline{B}G_{t-1}\right) \\ Pr\left(GG_t|BB_{t-1}\right) & Pr\left(\overline{G}B_t|BB_{t-1}\right) & Pr\left(\overline{B}G_t|BB_{t-1}\right) & Pr\left(BB_t|BB_{t-1}\right) \end{pmatrix}.$$

If the Markov chain characterised by T is irreducible it has a unique stationary distribution which is defined by the row vector (1.8) that satisfies both

$$\begin{aligned} \left(\rho^{GG}\left(\widehat{m}\right),\rho^{\overline{G}B}\left(\widehat{m}\right),\rho^{\overline{B}G}\left(\widehat{m}\right),\rho^{BB}\left(\widehat{m}\right)\right) &= \\ \left(\rho^{GG}\left(\widehat{m}\right),\rho^{\overline{G}B}\left(\widehat{m}\right),\rho^{\overline{B}G}\left(\widehat{m}\right),\rho^{BB}\left(\widehat{m}\right)\right)T & \end{aligned} \tag{1.12}$$

and

$$\rho^{GG}\left(\widehat{m}\right)+\rho^{\overline{G}B}\left(\widehat{m}\right)+\rho^{\overline{B}G}\left(\widehat{m}\right)+\rho^{BB}\left(\widehat{m}\right)=1. \tag{1.13}$$

Further, these conditions are sufficient for the transition probabilities to converge to the stationary distribution:

$$\lim_{k\to\infty} T^k = \mathbf{1}\left(\rho^{GG}\left(\widehat{m}\right),\rho^{\overline{G}B}\left(\widehat{m}\right),\rho^{\overline{B}G}\left(\widehat{m}\right),\rho^{BB}\left(\widehat{m}\right)\right),$$

where $\mathbf{1}$ is a 4×1 vector with all entries equal to 1.

Each transition probability is comprised of two parts: i) the undersubscribed losing school in $t-1$ is replaced with probability γ by a new school which is good with probability θ, and ii) the oversubscribed winning school in $t-1$ admits a share $\widehat{m}p$ of high-ability non-local students. The pair of school qualities and the student allocation then determine the winner in period t. The transition matrix T has the following entries:

$$\begin{aligned} Pr\left(GG_t|GG_{t-1}\right) &= \gamma\theta+(1-\gamma) \\ Pr\left(\overline{G}B_t|GG_{t-1}\right) &= \gamma(1-\theta)\frac{1}{2}\left(s+\widehat{m}pd\right) \\ Pr\left(\overline{B}G_t|GG_{t-1}\right) &= \gamma(1-\theta)\left[1-\frac{1}{2}\left(s+\widehat{m}pd\right)\right] \\ Pr\left(BB_t|GG_{t-1}\right) &= 0 \end{aligned}$$

$$
\begin{aligned}
Pr\left(GG_t|\overline{G}B_{t-1}\right) &= \gamma\theta \\
Pr\left(\overline{G}B_t|\overline{G}B_{t-1}\right) &= [1-\gamma+\gamma(1-\theta)]\frac{1}{2}(s+\widehat{m}pd) \\
Pr\left(\overline{B}G_t|\overline{G}B_{t-1}\right) &= [1-\gamma+\gamma(1-\theta)]\left[1-\frac{1}{2}(s+\widehat{m}pd)\right] \\
Pr\left(BB_t|\overline{G}B_{t-1}\right) &= 0
\end{aligned}
$$

$$
\begin{aligned}
Pr\left(GG_t|\overline{B}G_{t-1}\right) &= 0 \\
Pr\left(\overline{G}B_t|\overline{B}G_{t-1}\right) &= [1-\gamma+\gamma(\theta)]\frac{1}{2}(s-\widehat{m}pd) \\
Pr\left(\overline{B}G_t|\overline{B}G_{t-1}\right) &= [1-\gamma+\gamma(\theta)]\left[1-\frac{1}{2}(s-\widehat{m}pd)\right] \\
Pr\left(BB_t|\overline{B}G_{t-1}\right) &= \gamma(1-\theta)
\end{aligned}
$$

$$
\begin{aligned}
Pr\left(BB_t|GG_{t-1}\right) &= 0 \\
Pr\left(BB_t|\overline{G}B_{t-1}\right) &= \gamma\theta\frac{1}{2}(s-\widehat{m}pd) \\
Pr\left(BB_t|\overline{B}G_{t-1}\right) &= \gamma\theta\left[1-\frac{1}{2}(s-\widehat{m}pd)\right] \\
Pr\left(BB_t|BB_{t-1}\right) &= 1-\gamma+\gamma(1-\theta)
\end{aligned}
$$

If $\gamma > 0$ and $0 < \theta < 1$ then the Markov chain defined by T is both irreducible. Solving for (1.8) using (1.12) and (1.13) yields

$$\rho^{GG}(\widehat{m}) = \frac{(s-p\cdot d\cdot\widehat{m})\,\theta^2}{(2\theta-1)(s-1)+1-p\cdot d\cdot\widehat{m}} \tag{1.14}$$

$$\rho^{\overline{G}B}(\widehat{m}) = \frac{(s-p\cdot d\cdot\widehat{m})\,\theta(1-\theta)}{(2\theta-1)(s-1)+1-p\cdot d\cdot\widehat{m}} \tag{1.15}$$

$$\rho^{\overline{B}G}(\widehat{m}) = \frac{(2-s-p\cdot d\cdot\widehat{m})\,\theta(1-\theta)}{(2\theta-1)(s-1)+1-p\cdot d\cdot\widehat{m}} \tag{1.16}$$

$$\rho^{BB}(\widehat{m}) = \frac{(2-s-p\cdot d\cdot\widehat{m})\,(1-\theta)^2}{(2\theta-1)(s-1)+1-p\cdot d\cdot\widehat{m}}. \tag{1.17}$$

Substituting (1.15) and (1.16) into (1.11) yields

$$I(\widehat{m}) = \frac{2(1-\theta)\,\theta(s-1)}{(2\theta-1)(s-1)+1-p\cdot d\cdot\widehat{m}}.$$

If $\gamma = 0$, then the Markov chain which moves between realizations of $\overline{G}B$ and $\overline{B}G$ is both irreducible. Applying the same method yields

$$\rho^{GG}(\widehat{m}) = \theta^2 \tag{1.18}$$

$$\rho^{\overline{G}B}(\widehat{m}) = \frac{(s - p \cdot d \cdot \widehat{m})\,\theta\,(1-\theta)}{1 - p \cdot d \cdot \widehat{m}} \tag{1.19}$$

$$\rho^{\overline{B}G}(\widehat{m}) = \frac{(2 - s - p \cdot d \cdot \widehat{m})\,\theta\,(1-\theta)}{1 - p \cdot d \cdot \widehat{m}} \tag{1.20}$$

$$\rho^{BB}(\widehat{m}) = (1-\theta)^2, \tag{1.21}$$

and hence

$$I(\widehat{m}) \quad = \quad \frac{2(1-\theta)\,\theta\,(s-1)}{1 - p \cdot d \cdot \widehat{m}}.$$

If $\theta = 0$, or if $\theta = 1$, there are either two good or two bad schools and $I(\widehat{m}) = 0$.

Equilibrium steady state: An equilibrium level of mobility satisfies condition (1.4). Let $G(\widehat{m}) \equiv F(V \cdot I(\widehat{m}))$. G is weakly monotone increasing and $\widehat{m} \in [0,1]$. By Tarski's fixed point theorem, there exists an m^* such that $G(m^*) = m^*$.

1.11.2 Convergence to Equilibrium Steady State: Proposition 2

1.11.2.1 Exogenous School Quality

Recurrence equation - If $\gamma = 0$, then informativeness depends on the conjectured mobility of past generations if and only if one school is good and one is bad (see (1.18)-(1.21)). Using transition probabilities derived in the proof of Proposition 1 and substituting $\gamma = 0$ and $\widehat{m} = m_{t-1}$ yields

$$\begin{aligned} Pr\left(\overline{G}B_{t-1}\right) &= Pr\left(\overline{G}B_{t-1}|\overline{G}B_{t-2}\right) Pr\left(\overline{G}B_{t-2}\right) \\ &\quad + Pr\left(\overline{G}B_{t-1}|\overline{B}G_{t-2}\right) Pr\left(\overline{B}G_{t-2}\right) \\ &= \frac{s}{2}\left(Pr\left(\overline{G}B_{t-2}\right) + P\left(\overline{B}G_{t-2}\right)\right) \\ &\quad + \frac{1}{2} p d m_{t-1}\left(Pr\left(\overline{G}B_{t-2}\right) - Pr\left(\overline{B}G_{t-2}\right)\right) \end{aligned} \tag{1.22}$$

and

$$\begin{aligned} Pr\left(BG_{t-1}\right) &= Pr\left(\overline{B}G_{t-1}|\overline{G}B_{t-2}\right)Pr\left(\overline{G}B_{t-2}\right) \\ &\quad +Pr\left(\overline{B}G_{t-1}|BG_{t-2}\right)Pr\left(\overline{B}G_{t-2}\right) \\ &= \frac{2-s}{2}\left(Pr\left(\overline{G}B_{t-2}\right)+Pr\left(\overline{B}G_{t-2}\right)\right) \\ &\quad -\frac{1}{2}pdm_{t-1}\left(Pr\left(\overline{G}B_{t-2}\right)-Pr\left(\overline{B}G_{t-2}\right)\right). \end{aligned} \tag{1.23}$$

Subtracting (1.23) from (1.22) yields

$$I_t = \left[(s-1)2\theta\left(1-\theta\right)+pdF\left(VI_{t-1}\right)I_{t-1}\right] \equiv Z\left(I_{t-1}\right), \tag{1.24}$$

since i) if $\gamma = 0$ then

$$Pr\left(\overline{G}B_{t-2}\right)+Pr\left(\overline{B}G_{t-2}\right) = 2\theta\left(1-\theta\right),$$

and since ii) by Bayesian updating it must hold that

$$\begin{aligned} I_{t-1} &\equiv \mu\left(\overline{G}B|W_{t-2}\right)-\mu\left(\overline{B}G|W_{t-2}\right) \\ &= Pr\left(\overline{G}B_{t-2}\right)-Pr\left(\overline{B}G_{t-2}\right), \end{aligned}$$

and since iii) each generation behaves optimally given their level of informativeness it must hold that,

$$m_{t-1} = F\left(VI_{t-1}\right).$$

Convergence - The sequence I_t converges to the smallest equilibrium level of informativeness. An increasing sequence converges to its least upper bound. Show 1) that the sequence is increasing and then 2) that the smallest fixed point is its least upper bound.

1) The sequence I_t is increasing because $Z\left(I_t\right)$ increases in I_t. For any I_t, $I_t^{'}$ such that $I_t^{'} > I_t^{'}$,

$$\begin{aligned} Z\left(I_t^{'}\right)-Z\left(I_t\right) &= pd\left\{\left[F\left(V\cdot I_t^{'}\right)-F\left(V\cdot I_t\right)\right]I_t+F\left(V\cdot I_t^{'}\right)\left[I_t^{'}-I_t\right]\right\} \\ &\geq 0, \end{aligned}$$

since F is positive and increasing and $V > 0$ and $I_t \geq 0$.

Using $I_0 = 0$, (1.24), and $I \in [0,1]$, then

$$I_1 = Z(I_0) \geq I_0 = 0,$$

and due to Z being increasing

$$Z(Z(I_{t-1})) = Z(I_t) \geq I_t = Z(I_{t-1})$$

for all I_t.

2) The smallest upper bound of the sequence I_t is given by the smallest fixed point $\bar{I}$, where

$$\bar{I} \equiv inf\{I : Z(I) \leq I\}.$$

If $\bar{I}$ was not an upper bound then for some $I^{'} \leq \bar{I}$ it would be true that $Z\left(I^{'}\right) > Z\left(\bar{I}\right)$. But Z is increasing and $Z\left(\bar{I}\right) = \bar{I}$, so this is a contradiction. Further, if $\bar{I}$ was not the least upper bound then for some I^{**}, $I^{**} < \bar{I}$. An increasing sequence converges to its least upper bound. In addition, the limit of the sequence is a fixed point such that: $Z(I^{**}) = I^{**}$. But I assumed that $\bar{I}$ was the smallest fixed point of Z.

By (1.1),

$$m_t = F(V \cdot I_t)$$

and $V > 0$ and F is increasing. m_t increases monotonically with I_t. Hence, as I_t converges so does m_t. Further, the smallest equilibrium level of informativeness I^{1*} corresponds to the smallest equilibrium level of mobility m^{1*}.

1.11.2.2 Endogenous School Quality

In generation 0,

$$\begin{aligned} Pr(GG_0) &= \theta^2 \\ Pr(BB_0) &= (1-\theta)^2 \\ I_0 &= 0. \end{aligned}$$

Using transition probabilities from the proof of Proposition 1 and evaluating them at $\widehat{m} = m_{t-1}$ yields the following system of recurrence equations

$$\begin{aligned} I_t &= [\gamma(1-\theta)(s-1+dpm_{t-1})]P(GG_{t-1}) \\ &\quad + [(1-\gamma+\gamma(1-\theta))(s-1+dpm_{t-1})]\frac{I_{t-1}+1-P(GG_{t-1})-P(BB_{t-1})}{2} \\ &\quad + [(1-\gamma+\gamma(\theta))(s-1-dpm_{t-1})]\frac{1-P(GG_{t-1})-P(BB_{t-1})-I_{t-1}}{2} \\ &\quad + [\gamma\theta(s-1-dpm_{t-1})]P(BB_{t-1}) \end{aligned}$$

$$P(GG_t) = \left[\gamma\theta + (1-\gamma) - \frac{1}{2}\gamma\theta\right]P(GG_{t-1}) + \frac{1}{2}\gamma\theta + \frac{1}{2}\gamma\theta I_{t-1} - \frac{1}{2}\gamma\theta P(BB_{t-1})$$

$$\begin{aligned} P(BB_t) &= \frac{1}{2}\gamma(1-\theta) - \frac{1}{2}\gamma(1-\theta)I_{t-1} - \frac{1}{2}\gamma(1-\theta)P(GG_{t-1}) \\ &\quad + \left[1-\gamma+\gamma(1-\theta) - \frac{1}{2}\gamma(1-\theta)\right]P(BB_{t-1}) \end{aligned}$$

where

$$m_{t-1} = F(V \cdot I_{t-1}).$$

Convergence - I will use simulations to support the argument that mobility and informativeness converge to the smallest steady-state equilibrium levels. In particular, I will assume F is a Uniform distribution on $[0,\overline{c}]$, and pick values of $\overline{c}$ such that the equilibrium steady-state level of informativeness and mobility is not unique. The parameters which influence how informativeness evolves over time are θ, γ, s and $\frac{V}{\overline{c}}dp$, where γ influences the rate of convergence but not the limit. The following graphs show that starting from $I_0 = 0$, the level of informativeness converges to the smallest equilibrium steady-state level for various values of θ.

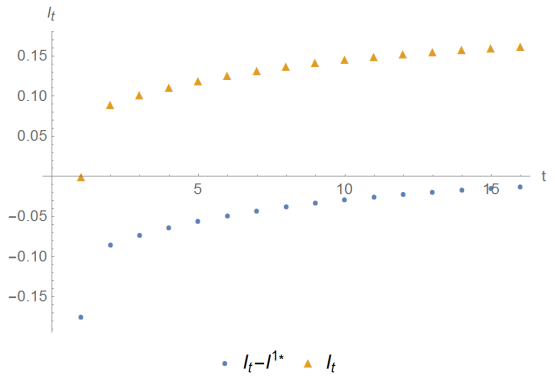

Figure 1.6: Parameter values: $\theta = \frac{1}{10}, \gamma = \frac{1}{2}, s = \frac{3}{2}, \frac{V}{c}dp = \frac{1}{2}$

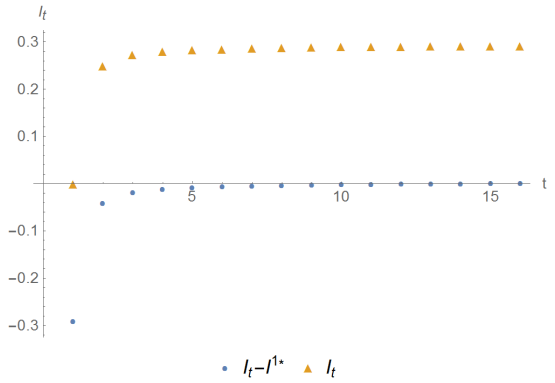

Figure 1.7: Parameter values: $\theta = \frac{1}{2}, \gamma = \frac{1}{2}, s = \frac{3}{2}, \frac{V}{c}dp = \frac{1}{2}$

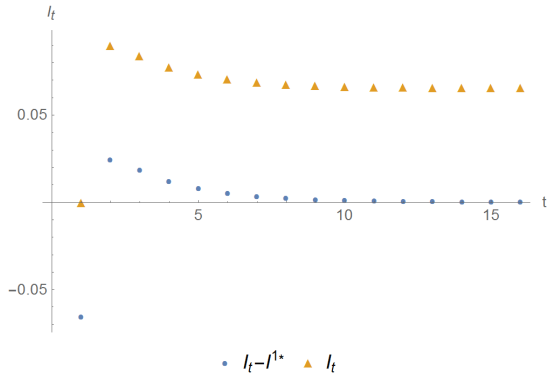

Figure 1.8: Parameter values: $\theta = \frac{9}{10}, \gamma = \frac{1}{2}, s = \frac{3}{2}, \frac{V}{c}dp = \frac{1}{2}$

1.11.3 Comparative Statics: Proposition 3

Let informativeness, I, be our variable of interest. Define a function

$$Z(I, \gamma, F(\cdot), p, d, s, \theta, V) \rightarrow [0, 1]$$

by

$$Z \equiv \frac{2(1-\theta)\theta(s-1)}{(2\theta-1)(s-1)1_{\gamma>0}(\gamma)+1-p\cdot d\cdot F(V\cdot I)}.$$

To explain the method used to derive comparative statics, denote the parameter subject to exogenous change by $r \in R$, where R is a partially ordered set. To simplify notation, I will write $Z(r,I) : T \times [0,1] \to [0,1]$. For any $r \in R$, $Z(r,I)$ is continuous but for upward jumps in I. Then I will show that $Z(r,I)$ is monotone non-decreasing in r for all $I \in [0,1]$. After Corollary 1, (p. 446), in Milgrom and Roberts [1994], this implies that the smallest fixed point of $Z(r,I)$, given by $I^{1*}(r)$, is monotone non-decreasing in r. Finally, I will show that $m^{1*}(r)$ is also monotone non-decreasing in r using condition (1.1).

1.11.3.1 Exogenous Quality

1. Since $Z(p,I)$ is non-decreasing in p at every I, I^{1*} is non-decreasing in p. Since $m^{1*} = F\left(V\cdot I^{1*}\right)$ and $V > 0$, m^{1*} also is non-decreasing in p.

2. Define $x(F(\cdot),I) \equiv F(V\cdot I)$. For any I, and any $F(\cdot)$ and $\overline{F}(\cdot)$, such that $F(\cdot)$ first-order stochastically dominates $\overline{F}(\cdot)$, it holds that

$$x(I,F(\cdot)) = F(V\cdot I) \leq \overline{F}(V\cdot I) = x\left(I,\overline{F}(\cdot)\right).$$

Since

$$Z(x(F(\cdot),I),I) = \frac{2(1-\theta)\theta(s-1)}{1-p\cdot d\cdot x(F(\cdot),I)}$$

is non-decreasing in x at every I, I^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$. Since $m^{1*} = F\left(V\cdot I^{1*}\right)$ and

$$F\left(V\cdot I^{1*}(F(\cdot))\right) \leq \overline{F}\left(V\cdot I^{1*}\left(\overline{F}(\cdot)\right)\right)$$

m^{1*} also is non-decreasing with a negative shift in FOSD of $F(\cdot)$.

1.11.3.2 Endogenous Quality

If $\gamma > 0$, the fraction of good schools in equilibrium steady state is given by

$$\Theta_{\gamma>0} = \rho^{GG}\left(m^{1*}\right) + \frac{1}{2}\rho^{\overline{G}B}\left(m^{1*}\right) + \frac{1}{2}\rho^{\overline{B}G}\left(m^{1*}\right) = \frac{\theta\left(\theta(s-1)+1-p\cdot d\cdot m^{1*}\right)}{(2\theta-1)(s-1)+1-p\cdot d\cdot m^{1*}}.$$

When analysing comparative statics on $\Theta_{\gamma>0}(r)$ with respect to any parameter r, I can decompose the total effect into a direct and an indirect effect:

$$\frac{d\Theta_{\gamma>0}}{dr} = \frac{\partial\Theta_{\gamma>0}}{\partial r} + \frac{\partial\Theta_{\gamma>0}}{\partial m^{1*}}\frac{\partial m^{1*}}{\partial r}, \tag{1.25}$$

where

$$\frac{\partial\Theta_{\gamma>0}}{\partial m^{1*}} = \frac{p\cdot d\cdot\theta(1-\theta)(s-1)}{\left[(2\theta-1)(s-1)+1-p\cdot d\cdot m^{1*}\right]^2} \geq 0. \tag{1.26}$$

1. Since $Z(p,I)$ is non-decreasing in p at every I, I^{1*} is non-decreasing in p. Since $m^{1*} = F\left(V\cdot I^{1*}\right)$ and $V>0$, m^{1*} also is non-decreasing in p. $\Theta_{\gamma>0}$ is non-decreasing in p since (1.25) is weakly positive given (1.26), and given that

$$\frac{\partial\Theta_{\gamma>0}}{\partial p} = \frac{m\cdot d\cdot\theta(1-\theta)(s-1)}{\left[(2\theta-1)(s-1)+1-p\cdot d\cdot m\right]^2} \geq 0$$

and given that m^{1*} is non-decreasing in p.

2. To simplify, define $x(F(\cdot),I) \equiv F(V\cdot I)$. For any I, and any $F(\cdot)$ and $\overline{F}(\cdot)$, such that $F(\cdot)$ first-order stochastically dominates $\overline{F}(\cdot)$, it holds that

$$x(I,F(\cdot)) = F(V\cdot I) \leq \overline{F}(V\cdot I) = x\left(I,\overline{F}(\cdot)\right).$$

Since

$$Z(x(F(\cdot),I),I) = \frac{2(1-\theta)\theta(s-1)}{(2\theta-1)(s-1)+1-p\cdot d\cdot x(F(\cdot),I)}$$

is non-decreasing in x at every I, I^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$. Since $m^{1*} = F\left(V\cdot I^{1*}\right)$ and

$$F\left(V\cdot I^{1*}(F(\cdot))\right) \leq \overline{F}\left(V\cdot I^{1*}\left(\overline{F}(\cdot)\right)\right),$$

m^{1*} also is non-decreasing with a negative shift in FOSD of $F(\cdot)$. $\Theta_{\gamma>0}$ is non-decreasing with a negative shift in FOSD of $F(\cdot)$ given (1.26) because $F(\cdot)$ affects $\Theta_{\gamma>0}$ only through its effect on m^{1*} and m^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$.

1.11.4 Comparative Statics: Proposition 4

For any I,

$$\begin{aligned} Z(\gamma>0,I)-Z(\gamma=0,I) &= \\ \frac{2(s-1)(1-\theta)\theta}{(2\theta-1)(s-1)+1-p\cdot d\cdot F(V\cdot I)}-\frac{2(s-1)(1-\theta)\theta}{1-p\cdot d\cdot F(V\cdot I)} &\geq 0 \end{aligned}$$

if and only if $\theta \leq \frac{1}{2}$. Hence, I^{1*} is non-decreasing as γ is raised above $\gamma=0$ if and only if $\theta \leq \frac{1}{2}$. Since $m^{1*}=F\left(V\cdot I^{1*}\right)$ and $V>0$, m^{1*} is non-decreasing as γ is raised above $\gamma=0$ if and only if $\theta \leq \frac{1}{2}$. The fraction of schools is raised since $\Theta_{\gamma=0}=\theta$ and for any m

$$\begin{aligned} \Theta_{\gamma>0}-\Theta_{\gamma=0} &= \frac{\theta(\theta(s-1)+1-p\cdot d\cdot m)}{(2\theta-1)(s-1)+1-p\cdot d\cdot m}-\theta \\ & \frac{(s-1)(1-\theta)\theta}{(2\theta-1)(s-1)+1-p\cdot d\cdot m}\geq 0. \end{aligned} \tag{1.27}$$

1.11.5 Welfare: Proposition 5

Denote the fraction of high-ability students at a good school in steady state by h^G, and the fraction of low-ability students at a good school in steady state by l^G. Since the share of high-ability students at the most recent winner is $\frac{1}{2}\left(1+m^{1*}p\right)$, and using the stationary probabilities derived in the proof of Proposition 1:

$$\begin{aligned} h^G &= \rho^{GG}\left(m^{1*}\right)+\frac{1}{2}\left(1+m^{1*}p\right)\rho^{\overline{G}B}\left(m^{1*}\right)+\frac{1}{2}\left(1-m^{1*}p\right)\rho^{\overline{B}G}\left(m^{1*}\right) \\ &= \frac{\theta\left(1+\theta(s-1)1_{\gamma>0}(\gamma)-m^{1*}p(d-(1-\theta)(s-1))\right)}{\left[(2\theta-1)(s-1)1_{\gamma>0}(\gamma)+1-dm^{1*}p\right]} \end{aligned}$$

$$\begin{aligned} l^G &= \rho^{GG}\left(m^{1*}\right)+\frac{1}{2}\left(1-m^{1*}p\right)\rho^{\overline{G}B}\left(m^{1*}\right)+\frac{1}{2}\left(1+m^{1*}p\right)\rho^{\overline{B}G}\left(m^{1*}\right) \\ &= \frac{\theta\left(1+\theta(s-1)1_{\gamma>0}(\gamma)-m^{1*}p(d+(1-\theta)(s-1))\right)}{\left[(2\theta-1)(s-1)1_{\gamma>0}(\gamma)+1-dm^{1*}p\right]}. \end{aligned}$$

When analysing comparative statics on $h^G(r)$ with respect to any parameter r, I can decompose the total effect into a direct and an indirect effect:

$$\frac{dh^G(r)}{dr}=\frac{\partial h^G(r)}{\partial r}+\frac{\partial h^G(r)}{\partial m^{1*}}\frac{\partial m^{1*}}{\partial r}, \tag{1.28}$$

where for $\gamma = 0$,

$$\frac{\partial h^G(r)}{\partial m^{1*}} = \frac{p(s-1)(1-\theta)\theta}{[1-dpm^{1*}]^2} \geq 0, \tag{1.29}$$

and for $\gamma > 0$,

$$\frac{\partial h^G(r)}{\partial m^{1*}} = \frac{p(s-1)(1-\theta)\theta((2\theta-1)(s-1)+1+d)}{[(2\theta-1)(s-1)+1-dpm^{1*}]^2} \geq 0. \tag{1.30}$$

Similarly,

$$\frac{dl^G(r)}{dr} = \frac{\partial l^G(r)}{\partial r} + \frac{\partial l^G(r)}{\partial m^{1*}} \frac{\partial m^{1*}}{\partial r}, \tag{1.31}$$

where for $\gamma = 0$,

$$\frac{\partial l^G(r)}{\partial m^{1*}} = -\frac{p(s-1)(1-\theta)\theta}{[1-dpm^{1*}]^2} \leq 0, \tag{1.32}$$

and where for $\gamma > 0$,

$$\frac{\partial l^G(r)}{\partial m^{1*}} = -\frac{p(s-1)(1-\theta)\theta((2\theta-1)(s-1)+1-d)}{[(2\theta-1)(s-1)+1-dpm^{1*}]^2} \leq 0. \tag{1.33}$$

1.11.5.1 Exogenous School Quality

1. h^G is non-decreasing in p given (1.29), since h^G is non-decreasing in p at any m^{1*}, and since m^{1*} is non-decreasing in p by Proposition 3.

2. $F(\cdot)$ affects h^G only through its effect on m^{1*}. Since m^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$ by Proposition 3, since (1.29) implies that h^G is non-decreasing with a negative shift in FOSD of $F(\cdot)$.

1.11.5.2 Endogenous School Quality

1. Since for any $m \in [0,1]$,

$$\frac{\partial h^G}{\partial p} = \frac{m(s-1)(1-\theta)\theta((2\theta-1)(s-1)+1+d)}{[(2\theta-1)(s-1)+1-dpm]^2} \geq 0,$$

and since m^{1*} is non-decreasing in p by Proposition 4, (1.30) implies that h^G is non-decreasing in p.

Since for any $m \in [0,1]$,

$$\frac{\partial l^G}{\partial p} = -\frac{m(s-1)(1-\theta)\theta((2\theta-1)(s-1)+1-d)}{[(2\theta-1)(s-1)+1-dpm]^2} \leq 0,$$

and since m^{1*} is non-decreasing in p by Proposition 4, (1.32) implies that l^G is non-increasing in p.

2. $F(\cdot)$ affects h^G only through its effect on m^{1*}. Since m^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$ by Proposition 4, (1.30) implies that h^G is non-decreasing with a negative shift in FOSD of $F(\cdot)$.

$F(\cdot)$ affects l^G only through its effect on m^{1*}. Since m^{1*} is non-decreasing with a negative shift in FOSD of $F(\cdot)$ by Proposition 4, (1.33) implies that l^G is non-increasing with a negative shift in FOSD of $F(\cdot)$.

1.11.6 Welfare: Proposition 6

Consider the impact on high-ability students. Decompose the total effect:

$$\begin{aligned} h^G_{\gamma>0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma=0}} &= h^G_{\gamma>0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma>0}} \\ &\quad + h^G_{\gamma=0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma=0}}. \end{aligned}$$

If $\theta \leq \frac{1}{2}$, then h^G weakly increases as γ is raised above 0.

$$h^G_{\gamma>0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma>0}} \geq 0 \tag{1.34}$$

holds, since the quality of schools improves as γ is raised above 0 at every m by (1.27), and since

$$h^G_{\gamma=0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma=0}} \geq 0$$

given $m^{1*}_{\gamma>0} \geq m^{1*}_{\gamma=0}$ by Proposition 4, and since (1.29) holds.

If $\frac{1}{2} < \theta \leq \frac{1-p(2+d-s)}{2p(s-1)}$, then although $m^{1*}_{\gamma>0} \leq m^{1*}_{\gamma=0}$ by Proposition 4, even if the drop in m^{1*} is as large as possible,

$$\begin{aligned} h^G_{\gamma>0,m^{1*}_{\gamma>0}=0} - h^G_{\gamma=0,m^{1*}_{\gamma=0}=1} &= \\ \frac{(s-1)(\theta-1)\theta(1-p(1+d+(s-1)(2\theta-1)))}{(1-dp)(1+(s-1)(2\theta-1))} &\geq 0. \end{aligned}$$

Therefore,

$$h^G_{\gamma>0,m^{1*}_{\gamma>0}} - h^G_{\gamma=0,m^{1*}_{\gamma=0}} \geq 0.$$

Consider the impact on low-ability students. Decompose the total effect:

$$\begin{aligned} l^G_{\gamma>0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma=0}} &= l^G_{\gamma>0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma>0}} \\ &\quad + l^G_{\gamma=0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma=0}}. \end{aligned}$$

If $\theta \geq \frac{1}{2}$, then l^G weakly increases as γ is raised above 0.

$$l^G_{\gamma>0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma>0}} \geq 0, \tag{1.35}$$

holds since the quality of schools improves as γ is raised above 0 at every m by (1.27), and because

$$l^G_{\gamma=0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma=0}} \geq 0$$

given $m^{1*}_{\gamma>0} \leq m^{1*}_{\gamma=0}$ by Proposition 4, and because (1.32) holds.

If $\theta \leq \frac{1}{2}$, then l^G weakly decreases as γ is raised above 0, because although $m^{1*}_{\gamma>0} \geq m^{1*}_{\gamma=0}$ by Proposition 4, even if the increase in m^{1*} is as large as possible then

$$\begin{aligned} l^G_{\gamma>0,m^{1*}_{\gamma>0}=1} - l^G_{\gamma=0,m^{1*}_{\gamma=0}=0} &= \\ \frac{(s-1)(1-\theta)\theta(1-p)}{(1-dp)+(s-1)(2\theta-1)} &\geq 0. \end{aligned}$$

Therefore,

$$l^G_{\gamma>0,m^{1*}_{\gamma>0}} - l^G_{\gamma=0,m^{1*}_{\gamma=0}} \geq 0.$$

1.11.7 Extensions: Longer Window of Performance-based Rankings

1.11.7.1 Exogenous School Qualities

Because school qualities are exogenous, I can focus on the case in which one school is good and one is bad. Due to the symmetry of the setting, I can restrict attention to distinguishing between two sets of histories observed by families in generation t: those in which the same school i won, denoted by H^{ii}_t, and those in which school i won in period $t-2$ and then school j won in period $t-1$, denoted by H^{ij}_t. Assume that, in each generation, conditional on history H^{ii} a share $\widehat{m}^{H^{ii}}$ of families travel to school i even though this is their non-local school, and assume that conditional on history H^{ij} a share $\widehat{m}^{H^{ij}}$ of families travel to school j even though this is their non-local school. Denote by $k_{t-2}l_{t-1}$ the event that the school

of quality k won in period $t-2$, and the school of quality l won in period $t-1$, where $k,l \in \{G,B\}$. Then the transition matrix is given by

$$T = \begin{pmatrix} Pr(G_{t-1}G_t|G_{t-2}G_{t-1}) & Pr(G_{t-1}B_t|G_{t-2}G_{t-1}) & Pr(B_{t-1}G_t|G_{t-2}G_{t-1}) & Pr(B_{t-1}B_t|G_{t-2}G_{t-1}) \\ Pr(G_{t-1}G_t|G_{t-2}B_{t-1}) & Pr(G_{t-1}B_t|G_{t-2}B_{t-1}) & Pr(B_{t-1}G_t|G_{t-2}B_{t-1}) & Pr(B_{t-1}B_t|G_{t-2}B_{t-1}) \\ Pr(G_{t-1}G_t|B_{t-2}G_{t-1}) & Pr(G_{t-1}B_t|B_{t-2}G_{t-1}) & Pr(B_{t-1}G_t|B_{t-2}G_{t-1}) & Pr(B_{t-1}B_t|B_{t-2}G_{t-1}) \\ Pr(G_{t-1}G_t|B_{t-2}B_{t-1}) & Pr(G_{t-1}B_t|B_{t-2}B_{t-1}) & Pr(B_{t-1}G_t|B_{t-2}B_{t-1}) & Pr(B_{t-1}B_t|B_{t-2}B_{t-1}) \end{pmatrix},$$

where

$$\begin{aligned} Pr(G_{t-1}G_t|G_{t-2}G_{t-1}) &= \frac{1}{2}\left(s+\widehat{m}^{H^{ii}}pd\right) \\ Pr(G_{t-1}G_t|B_{t-2}G_{t-1}) &= \frac{1}{2}\left(s+\widehat{m}^{H^{ij}}pd\right) \end{aligned}$$

$$\begin{aligned} Pr(G_{t-1}B_t|G_{t-2}G_{t-1}) &= 1-\frac{1}{2}\left(s+\widehat{m}^{H^{ii}}pd\right) \\ Pr(G_{t-1}B_t|B_{t-2}G_{t-1}) &= 1-\frac{1}{2}\left(s+\widehat{m}^{H^{ij}}pd\right) \end{aligned}$$

$$\begin{aligned} Pr(B_{t-1}G_t|G_{t-2}B_{t-1}) &= \frac{1}{2}\left(s-\widehat{m}^{H^{ij}}pd\right) \\ Pr(B_{t-1}G_t|B_{t-2}B_{t-1}) &= \frac{1}{2}\left(s-\widehat{m}^{H^{ii}}pd\right) \end{aligned}$$

$$\begin{aligned} Pr(B_{t-1}B_t|G_{t-2}B_{t-1}) &= 1-\frac{1}{2}\left(s-\widehat{m}^{H^{ij}}pd\right) \\ Pr(B_{t-1}B_t|B_{t-2}B_{t-1}) &= 1-\frac{1}{2}\left(s-\widehat{m}^{H^{ii}}pd\right) \end{aligned}$$

and all other transitions have zero probability.

Denote by $Pr(kl)$ the steady-state probability of the event that the school of quality k won in a given period and the school of quality l won in the subsequent period, where $k,l \in \{G,B\}$. These probabilities are derived using the same methodology as in the proof of Proposition 1. Given families use Bayes' rule to update their beliefs, informativeness (i.e. likelihood that the most recent winner is good and the most recent loser is bad minus the probability that the reverse is true) conditional on observing history H^{ij} is given by

$$I^{H^{ij}} = Pr(BG) - Pr(GB) = 0$$

for any values of $\widehat{m}^{H^{ii}}$ and $\widehat{m}^{H^{ij}}$. Further, informativeness conditional on observing history H^{ii} is given by

$$I^{H^{ii}} = Pr(GG) - Pr(BB) = \frac{s-1}{1-dp\widehat{m}^{H^{ii}}} \geq 0$$

for any values of $\widehat{m}^{H^{ii}}$ and $\widehat{m}^{H^{ij}}$.

If $\widehat{m}^{H^{ij}}$ increases, then the levels of informativeness $I^{H^{ij}}$ and $I^{H^{ii}}$ remain unchanged. If $\widehat{m}^{H^{ii}}$ increases, $I^{H^{ij}}$ remains unchanged, while $I^{H^{ii}}$ weakly increases. Optimal mobility conditional on history H^{ii} is given by $m^{H^{ii}} = F\left(V \cdot I^{H^{ii}}\right)$, and optimal mobility conditional on history H^{ij} is given by $m^{H^{ij}} = F\left(V \cdot I^{H^{ij}}\right)$. Hence optimal mobility is nondecreasing in conjectured levels of mobility $\widehat{m}^{H^{ij}}$ and $\widehat{m}^{H^{ii}}$.

1.11.7.2 Endogenous School Qualities

Families still only distinguish between two possible kinds of histories: H^{ii} and H^{ij}. Define $\widehat{m}^{H^{ij}}$ and $\widehat{m}^{H^{ii}}$ as in the case of exogenous school qualities. There are now 16 possible states in which the process can be in period t, each state specifying the pair of school qualities in period $t-1$ and in period $t-2$ and the winning school in of these two periods. Denote a state in period t by $\left(\overline{k}l,\overline{r}t\right)_{ii}$ if the same school won in both periods and by $\left(\overline{k}l,\overline{r}t\right)_{ij}$ if a different school won in each period, where k is the quality of the winning school in period $t-2$, and l is the quality of the losing school in period $t-2$, r is the quality of the winning school in period $t-1$ and t is the quality of the losing school in period where $k,l,r,t \in \{G,B\}$. Transitions between states now depend on which school wins, as well as on whether the loser is replaced by a new school or not.

Given the stationary distribution, and given that families use Bayes' rule to update their posterior beliefs, informativeness conditional on observing history H^{ij} is given by

$$\begin{aligned} I^{H^{ij}} &= Pr\left(\left(\overline{B}B,\overline{G}B\right)_{ij}\right) + Pr\left(\left(\overline{B}G,\overline{G}B\right)_{ij}\right) \\ &\quad -Pr\left(\left(\overline{G}G,\overline{B}G\right)_{ij}\right) - Pr\left(\left(\overline{G}B,\overline{B}G\right)_{ij}\right). \end{aligned}$$

Informativeness conditional on observing history H^{ii} is given by

$$\begin{aligned} I^{H^{ii}} &= Pr\left(\left(\overline{G}G,\overline{G}B\right)_{ii}\right) + Pr\left(\left(\overline{G}B,\overline{G}B\right)_{ii}\right) \\ &\quad -Pr\left(\left(\overline{B}B,\overline{B}G\right)_{ii}\right) - Pr\left(\left(\overline{B}G,\overline{B}G\right)_{ii}\right). \end{aligned}$$

The following results are based on simulations in Mathematica and the code is available upon request: $I^{H^{ij}} \approx 0$ for all values of $\widehat{m}^{H^{ij}}$ and $\widehat{m}^{H^{ii}}$, and $I^{H^{ij}}$ increases in both $\widehat{m}^{H^{ij}}$ and $\widehat{m}^{H^{ii}}$. These simulations suggest that optimal mobility is increasing in conjectured mobility even if school qualities are endogenous.

1.11.8 Extensions: No Ability-based Selection by Schools

Assume families with a high-ability child derive a higher payoff from a good school than families with a low-ability child, i.e. $V_H > V_L$, and families observe the ability of their own child. Assume that an oversubscribed school uses a lottery to assign places to applicants, and that the chances of being allocated a place are independent of location. For simplicity, assume that school qualities are exogenously determined and that one of the schools is good and one is bad, but families do not observe which is which.

Let the mobility level characterizing families' optimal strategy conditional on the ability of their child be denoted by m_H and m_L respectively, where

$$m_H = F\left(V_H \cdot I\right)$$

$$m_L = F\left(V_L \cdot I\right).$$

The share of high-ability students applying to the better-performing school is $\frac{1}{2} + \frac{1}{2}m_H$ and the share of low-ability students applying is $\frac{1}{2} + \frac{1}{2}m_L$. Therefore, the share of the population of high-ability students admitted by the better-performing school is given by

$$\frac{1+m_H}{2+m_H+m_L}.$$

Applying the same methodology as in the proof of Proposition 1: Given transition matrix T

$$\begin{aligned} T &= \begin{pmatrix} Pr\left(GB_t|GB_{t-1}\right) & Pr\left(BG_t|GB_{t-1}\right) \\ Pr\left(GB_t|BG_{t-1}\right) & Pr\left(BG_t|BG_{t-1}\right) \end{pmatrix} \\ &= \begin{pmatrix} \frac{s}{2} + \left(\frac{m_H - m_L}{2+m_H+m_L}\right) d & 1 - \frac{s}{2} - \left(\frac{m_H - m_L}{2+m_H+m_L}\right) d \\ \frac{s}{2} + \left(\frac{m_L - m_H}{2+m_H+m_L}\right) d & 1 - \frac{s}{2} - \left(\frac{m_L - m_H}{2+m_H+m_L}\right) d \end{pmatrix}. \end{aligned}$$

Bayesian updating yields informativeness:

$$I(m_H, m_L) = \frac{(s-1)(2+m_H+m_L)}{2+m_H+m_L-2d(m_H-m_L)}.$$

Note that if $m_H = m_L$ then $I = 1-s$ which is the level of informativeness if students were randomly allocated to schools.

Define

$$Z(I) \equiv \frac{(s-1)(2+F(V_H \cdot I)+F(V_L \cdot I))}{2+F(V_H \cdot I)+F(V_L \cdot I)-2d(F(V_H \cdot I)-F(V_L \cdot I))}. \tag{1.36}$$

By assumption on the payoff,

$$F(V_H \cdot I) > F(V_L \cdot I). \tag{1.37}$$

Given (1.37), $Z(I)$ is non-decreasing in I if the distribution of transport costs is such that for all I,

$$\begin{aligned} \frac{\partial F(V_H \cdot I)}{\partial I} &\geq \frac{\partial F(V_L \cdot I)}{\partial I} \\ F'(V_H \cdot I) V_H &\geq F'(V_L \cdot I) V_L. \end{aligned} \tag{1.38}$$

Assuming (1.38) holds. Then by Tarski's fixed point theorem, a fixed point satisfying $I = Z(I)$ exists. Applying the same methodology as in the proof of the Proposition 3: Exogenous changes in the distribution of transport costs which cause $F(V_H \cdot I)$ to increase proportionally more than $F(V_L \cdot I)$ at any given level of I, increase $Z(I)$ at every I. To show this, substitute $x = F(V_L \cdot I)$ and $ax = F(V_H \cdot I)$ into (1.36) where $a > 1$, $1 > x > 0$ and $1 \geq xa > 0$, then

$$\frac{\partial}{\partial x} Z(I) \equiv \frac{\partial}{\partial x} \frac{(s-1)(2+ax+x)}{2+ax+x-2dx(a-1)} \geq 0.$$

Hence, the smallest equilibrium level of informativeness and the smallest levels of mobility for each student type are non-decreasing in such changes in the distribution F.

1.11.9 Extensions: Peer Effects

Assume that a family with transport cost c whose local school is school i derives the expected payoff

$$V \cdot Pr(W_t = i)$$

from attending school i and

$$V \cdot Pr(W_t = j) - c$$

from attending school j.

For simplicity, assume that schools' qualities are exogenously determined and that one of the schools is good and one is bad, but families do not know which is which. Generation t's optimal cut-off level for transport costs is the expected increase in the probability of attending the school which will win in the current period (W_t) if they apply to the school which won in the previous period (W_{t-1}):

$$\begin{aligned} C_t &= V \cdot \left[\mu\left(\overline{G}B|W_{t-1}\right)(s-1+\widetilde{m}pd) - \mu\left(\overline{B}G|W_{t-1}\right)(s-1-\widetilde{m}pd)\right] \\ &= V \cdot \left[(s-1)I_t + \widetilde{m}pd\right], \end{aligned}$$

where $\widetilde{m}$ is the current generation's level of mobility. Applying to the non-local losing school is no longer a dominated strategy. Let negative values of mobility represent the share of families who apply to the non-local losing school, and consistent with this allow F to also evaluate negative values such that $F(-c) = -F(c)$. For application strategies to be optimal, mobility $m \in [-1,1]$ needs to satisfy

$$m_t = F\left[V \cdot C_t\right]. \tag{1.39}$$

In equilibrium, families' beliefs about the mobility levels of past generations and the current generation need to be consistent with equilibrium. Substituting for steady-state informativeness (1.3) for the case of one good and one bad school with fixed qualities into (1.39), an equilibrium level of mobility must satisfy

$$m^* = F\left[V \cdot \left(\frac{(s-1)^2}{1-p\cdot d\cdot m^*} + m^* pd\right)\right]$$

and the corresponding equilibrium level of informativeness is given by

$$I^* = \frac{s-1}{1-p\cdot d\cdot m^*}.$$

Define

$$Z(m) \equiv F\left[V \cdot \frac{s-1}{1-p\cdot d\cdot m}\left[s-1+p\cdot d\cdot m\right]\right].$$

Since $Z(m)$ is non-decreasing in m and $m \in [-1,1]$, a fixed point of m such that $m = Z(m)$ exists by

Tarski's fixed point theorem.

The smallest (positive) equilibrium levels of mobility and informativeness increase as i) peer effects are stronger (d larger), or as ii) an oversubscribed school selects a larger proportion p of non-local high-ability applicants.

Applying the same methodology as in the proof of the Proposition 3:

- For any $m \in [0,1]$, $Z(m)$ is non-decreasing in p. Further, $Z(m)$ is non-decreasing in m. Therefore, the smallest fixed point $m \in [0,1]$, such that $Z(m) = m$ is non-decreasing in p.
- For any $m \in [0,1]$, $Z(m)$ is non-decreasing in d. Further, $Z(m)$ is non-decreasing in m. Therefore, the smallest fixed point of $m \in [0,1]$, such that $Z(m) = m$ is non-decreasing in d.

In the presence of peer effects, some share of families may apply to the non-local losing school, i.e. a negative equilibrium level of mobility may exist, and it is more likely to exist as i) peer effects are stronger (d larger), or as ii) an oversubscribed school selects a larger proportion p of non-local high-ability applicants.

A negative equilibrium level of mobility exists if and only if $m = Z(m)$ for some $m < 0$. Since there exist feasible parameter combinations such that $V \cdot \left[(s-1)\frac{s-1}{1-p\cdot d\cdot m} + p \cdot d \cdot m\right] < 0$ for $m < 0$, there exists some F such that $m = Z(m)$ for some $m < 0$. Such a fixed point is weakly more likely to exist the lower $V \cdot \left[(s-1)\frac{s-1}{1-p\cdot d\cdot m} + p \cdot d \cdot m\right]$ at any $m < 0$. Therefore, it is more likely to exist if d increases or if p increases.

1.11.10 Extensions: Signal of School Quality

Assume noise is additive. Let performance of a school of quality j with share of high-ability students h be given by $f(j,h) + \varepsilon_j$ where ε_j is an error term. Further, assume $\varepsilon_j - \varepsilon_i$ is distributed symmetrically about 0 with cdf $G(\cdot)$ and density $g(\cdot)$. Assume that $g(\cdot)$ is unimodal and is continuously differentiable and strictly positive everywhere. For the equilibrium levels of mobility and informativeness to increase as an oversubscribed school selects a larger share of non-local high-ability applicants or as lower transport costs are more likely (in a negative FOSD sense), it is sufficient that $f(G,h)$ and $f(B,h)$ satisfy the

following properties for very h:

$$\begin{aligned} Pr\left(f\left(G,h\right)-f\left(B,1-h\right)+\varepsilon_G-\varepsilon_B>0\right) &> Pr\left(f\left(B,h\right)-f\left(G,1-h\right)+\varepsilon_B-\varepsilon_G>0\right) \\ f\left(G,h\right)-f\left(B,1-h\right) &> f\left(B,h\right)-f\left(G,1-h\right) \end{aligned} \tag{1.40}$$

which corresponds to a performance advantage derived from better quality, and

$$\begin{aligned} \frac{\partial}{\partial h}Pr\left(f\left(G,h\right)-f\left(B,1-h\right)+\varepsilon_G-\varepsilon_B>0\right) &\geq \frac{\partial}{\partial h}Pr\left(f\left(B,h\right)-f\left(G,1-h\right)+\varepsilon_B-\varepsilon_G>0\right) \\ \frac{g\left(f\left(G,h\right)-f\left(B,1-h\right)\right)}{g\left(f\left(B,h\right)-f\left(G,1-h\right)\right)} &\geq \frac{\frac{\partial}{\partial h}\left(f\left(B,h\right)-f\left(G,1-h\right)\right)}{\frac{\partial}{\partial h}\left(f\left(G,h\right)-f\left(B,1-h\right)\right)}. \end{aligned} \tag{1.41}$$

By symmetry of $g\left(\cdot\right)$ and by (1.40), it must be true that

$$g\left(f\left(G,h\right)-f\left(B,1-h\right)\right)<g\left(f\left(B,h\right)-f\left(G,1-h\right)\right).$$

So for (1.41) to hold, complementarities in performance between student ability and school quality need to be sufficiently strong, i.e.

$$\frac{\frac{\partial}{\partial h}\left(f\left(G,h\right)-f\left(B,1-h\right)\right)}{\frac{\partial}{\partial h}\left(f\left(B,h\right)-f\left(G,1-h\right)\right)}$$

needs to be sufficiently large.

Chapter 2

Hidden Testing and Selective Disclosure of Evidence

Abstract

A decision maker (DM) can consult an advisor before choosing whether or not to switch from the status quo. Both agree that switching is optimal if and only if a certain hypothesis is true. But the advisor may be biased in how he trades off falsely accepting against falsely rejecting the hypothesis. The advisor has access to a testing technology which generates noisy but informative binary signals. Over two periods, the advisor can sequentially carry out costly tests before reporting his findings to the DM. I contrast the setting in which the DM observes all test outcomes with a setting in which testing itself is hidden and the advisor verifiably discloses outcomes. I fully characterise the circumstances in which the DM is strictly better off with hidden testing than with observable testing. Depending on how strongly the DM is inclined to stick with the status quo, hidden testing can be beneficial either if the advisor is more enthusiastic or if he is more reluctant to switch than the DM. Finally, I study what my findings imply about the bias of the ideal advisor.

2.1 Introduction

A decision maker faced with a choice under uncertainty will often have to rely on an advisor to gather evidence. In many cases, the advisor has incentives to acquire evidence because he cares about the decision being taken. However, the advisor may be biased relative to the DM regarding which decision he considers appropriate given the evidence. For example, a regulator relies on a company to produce evidence about the safety of a new product, but the regulator may be relatively more hesitant about licensing an unsafe product than about barring a safe product. Or a legislator relies on a scientist to produce research about the consequences of a policy, but the legislator may be relatively more anxious to introduce an effective measure than to prevent an ineffective measure.

The advisor often has to back up his recommendation with hard evidence, e.g. the company has to show proof of the test results, the scientist has to reveal his experimental data. In addition, the advisor may be restricted in how he acquires this evidence, e.g. the company may have to run a certified safety test, the scientist has to apply a state-of-the-art methodology. Despite these constraints, the advisor can strategically influence the decision maker, because he can choose *based on which findings to stop or to continue gathering evidence,* up to the point at which he must report his findings to the decision maker. Which degree of transparency is most effective in limiting the scope for manipulation? Should the decision maker require the advisor to disclose all evidence he has collected or should the decision maker allow the advisor to gather evidence in private and strategically choose which pieces of evidence to reveal?

In the model of this paper, a decision maker (she) can consult an advisor (he) before taking a binary decision. Both parties are uncertain whether or not a hypothesis is true. Both agree that it would be best to switch if the hypothesis is true and best to stick with the status quo if the hypothesis is false. The decision-maker's preferences characterise her predisposition to switch. The advisor may trade off the mistake of falsely accepting and the mistake of falsely rejecting differently than the decision maker (DM) and hence he may have a different predisposition to switch than the DM. I will refer to the action that the DM would choose given the common prior belief as "the default action". The advisor has access to a test technology which generates a noisy but informative binary signal at a fixed cost. In the first period, the advisor can choose whether or not to conduct a test and he will learn the outcome (signal realisation) at the end of the period. In the second period, he can choose whether or not to conduct an additional test. At the end of the second period, the DM receives the advisor's report and chooses

whether or not to switch. I contrast a setting with full transparency about testing to a setting with limited transparency. When testing is observable, the DM observes all test outcomes. When testing is hidden, the advisor tests in private and selectively discloses findings. Crucially, the DM does not observe at what point the evidence presented was generated or how much evidence was generated overall.

My main result (Theorem 1) fully characterises the situations in which the DM is better off if she allows the advisor to collect evidence in private rather than publicly. This possibility may seem counter-intuitive at first, because a biased advisor has even more scope for manipulation when he can hide test outcomes. However, the opportunity to selectively disclose outcomes alters the advisor's incentives for testing. When the advisor's incentives are weakly increased, then the advantageous effect of more testing may outweigh the disadvantageous effect of omitted evidence. In one situation the advisor cares relatively more about accepting the hypothesis if it is true, i.e. he is more enthusiastic about switching than the DM. In the other situation, the advisor cares relatively more about rejecting the hypothesis if it is false, i.e. he is more reluctant than the DM.

First, consider the situation in which the advisor is more enthusiastic than the DM. As far as the DM's preferences are concerned, if two tests were conducted then she would switch if and only if both outcomes are positive, and if one test was conducted then she would switch if and only if the outcome was positive. By contrast, if the advisor were to decide he would want the DM to switch from her default action for weaker evidence. Under observable testing, the advisor would stop testing after the first period. If the test outcome is positive, then the DM will switch action, in line with the advisor's interest. If the test outcome is negative, then the advisor may dislike that the DM will stick with her default action, but he has no hope of convincing her by conducting an additional test. Under hidden testing, the advisor will not stop testing after the first period, but instead conduct an additional test following a positive outcome. Although observing a single positive outcome is sufficient for the DM to switch according to her preferences, the advisor claiming that there is a single positive outcome is insufficient. This is because there is a risk that the advisor has omitted contradicting evidence, i.e. he may have tested twice and found a negative outcome on the first test. In order to convince the DM to switch, the advisor is forced to continue testing following a positive outcome in order to have a chance of presenting two positive outcomes. Therefore, since the advisor has an incentive to produce stronger evidence, the DM is better off.

Second, consider the situation in which the advisor is more reluctant than the DM. Suppose preferences are such that if two tests were conducted then i) the DM would switch even if one outcome was

positive and one was negative, while ii) the advisor would wish to switch if and only if both tests yielded a positive outcome.[1] Under observable testing, the advisor is deterred from conducting any tests. If the outcome of the first test was negative, the DM would not deviate from her default action. If the outcomes was positive the DM will be convinced to switch, contrary to the advisor's preferences, and the advisor cannot cause her to revert back to her default action by presenting further evidence. However, if testing is hidden then the advisor does have an incentive to test because he can ensure that the DM switches if and only if he himself believes that this is the appropriate decision. This is because the advisor can continue to gather evidence until he himself can be persuaded that switching is appropriate, but hide positive evidence if he himself is not convinced. Hidden testing thus provides the advisor with incentives to test, thereby benefiting the DM.

In both situations, the advisor has stronger incentives to test when testing is hidden, but the reasons why are fundamentally different.[2] When the advisor is more enthusiastic about switching than the DM and testing is hidden, then the advisor is more motivated to test because the DM requires stronger evidence to be convinced than under observable testing. The advisor can no longer persuade the DM with a single positive outcome because the DM is skeptical that this is not the only outcome the advisor has observed. By contrast, when the advisor is more reluctant than the DM and testing is hidden, the advisor is more motivated to test because he has greater influence over how the DM will act on the observed test outcomes. The advisor can explore whether or not he himself would be convinced to switch without fearing any adverse consequences for the DM's choice in case he discovers only weak evidence in favour of the hypothesis, because he then still has the option to withhold this evidence and the DM will stick with her default action.

Since the advisor's preferences affect his motivation to collect evidence, it is a natural question to ask which type of advisor the DM should choose to consult. I show that choosing an unbiased advisor is not always the best option and that either more enthusiastic or more reluctant advisors may be preferred, depending on the circumstances.

When choosing the ideal advisor, the DM should take into account how much the advisor's preferences diverge from her own, how costly testing is for the advisor and whether testing is hidden or observable. When testing is observable, the DM faces the following trade-off when choosing the ad-

[1] Assume that the test technology is such that a false negative outcome is more likely to arise than a false positive outcome.

[2] Proposition 1 highlights this contrast by comparing the equilibrium outcomes with those arising if the DM delegates the decision making to the advisor or if she has the power to commit to how she acts based on the evidence presented.

visor's bias: An unbiased advisor is motivated to present the DM with valuable evidence by following the same testing procedure the DM would follow. However, since testing is costly, he may not find the test worthwhile. By contrast, a biased advisor may not agree with the DM about what action is appropriate given the existing evidence, but he may be motivated to gather additional evidence to bring the DM's choice in line with what he himself would choose.

The DM values an advisor who collects evidence in such a way that his findings can be pivotal to which action she takes. Suppose the cost is high enough such that an unbiased advisor does not find it worthwhile to conduct enough tests to challenge the DM's default choice. Then she is better off with an advisor who is more enthusiastic about switching and therefore willing to incur higher costs to persuade her to switch. Next, suppose the cost is low enough such that an unbiased advisor produces some evidence which induces the DM to switch but the cost is too high for him to back up this evidence with an additional test. Then the DM is better off with an advisor who is more reluctant towards switching because this advisor is more willing to incur costs to check whether the earlier evidence is robust (Proposition 3). These two results hold irrespective of whether or not testing is observable.

By contrast, if testing is hidden there are further aspects to consider when choosing an advisor. An unbiased advisor has no incentive to hide test outcomes, which is per se desirable for the DM. However, as mentioned before, hidden testing alters a biased advisor's incentives for testing, which can benefit the DM depending on the divergence of preferences. Under hidden testing, the DM's preferences over the advisor's bias are not necessarily single peaked (Proposition 4). Suppose again the cost is low enough such that an unbiased advisor produces some evidence which induces the DM to switch but the cost is too high for him to back up this evidence with an additional test. Then not only a more reluctant but also a more enthusiastic advisor may conduct an additional test to validate earlier positive findings. This is because under hidden testing the DM finds a report of weak positive evidence less convincing when presented by a more enthusiastic advisor. She suspects that this advisor may be hiding negative evidence. In order to overcome the DM's skepticism a more enthusiastic advisor has to show that the earlier positive findings are backed up by an additional test.

My work builds on the literature studying verifiable disclosure when the collection of evidence is endogenous. Unlike other papers in the literature, I focus on the impact of hidden versus observable testing when the advisor's preferences over actions are state-dependent. Henry [2009] consider the impact of hidden testing when the advisor prefers to switch independent of the state (more enthusiastic) and has

to commit to the number of signals ex ante. I find that hidden testing is also beneficial when the advisor is more reluctant than the DM, a result that is independent of whether testing is simultaneous or sequential. Felgenhauer and Schulte [2014] study a persuasion game with an infinite time horizon in which testing is always hidden. They find that the lower the advisor's cost of testing the stronger is the evidence required by the DM to change from her default action. The advisor in their model is a special case of the advisor in my model and, unlike Felgenhauer and Schulte [2014], I highlight the implication of hidden versus observable testing for this advisor. Che and Kartik [2009] study a DM's preferences over the advisor's bias when the advisor can privately invest in the chance of obtaining a single signal and the action and state space are continuous. They also find that a more biased advisor can sometimes be preferred to an unbiased advisor, but their set-up is such that the direction of the bias plays no role. Finally, Gerardi and Yariv [2008] study the optimal mechanism to elicit information for a DM who faces a binary choice, when the DM can consult an arbitrary number of advisors and can also choose the bias of these advisors. In their model, the optimal bias is always such that the advisor is more enthusiastic about switching than the DM, whereas I find that the optimal bias may also be such that the advisor is more reluctant to switch than the DM. A more detailed discussion of the related literature can be found after the analysis is completed.

The paper is structured as follows. Section 2.2 introduces the model. Section 2.3 analyses a benchmark case in which a single player collects evidence and then takes a decision. The following sections study the model in which a DM consults an advisor. Section 2.4 analyses equilibrium outcomes when testing is observable. Section 2.5 analyses equilibrium outcomes when testing is hidden and establishes the main result. Section 2.6 analyses which bias the DM would prefer the advisor to have. Section 2.7 gives a more detailed discussion of the related literature. Section 2.8 concludes. All proofs can be found in the Appendix in Section 2.9.

2.2 Model

A decision maker (DM) must take a binary action and her loss from the action depends on the state of the world. The DM is not able to gather information or finds it prohibitively costly to do so, but she can consult an advisor (A). The advisor has a fixed time horizon in which he can test sequentially at a private cost before reporting to the DM.

2.2.1 General Framework

There are two states of the world, $\omega \in \{\omega^T, \omega^F\}$, which correspond to a novel hypothesis being either true (ω^T) or false (ω^F). The DM can choose an action $a \in \{a_R, a_S\}$, which is either to remain (a_R) or to switch (a_S). The common prior belief that the hypothesis is true is denoted by q. As shown in the next paragraph, the DM's preference will be chosen such that her default choice is a_R.

Preferences - A player i incurs a state-dependent loss $L_i(a, \omega)$ from the DM's chosen action and the loss is represented as follows:

$L_i(a,\omega)$	ω^F	ω^T
a_R	0	x_i
a_S	y_i	0

where $x_i, y_i \geq 0$ and $i = DM, A$. Hence, if the state of the world were known then players would agree on what the optimal action is, i.e. each player would prefer a_S if and only if $\omega = \omega^T$. However, each player may evaluate a loss from an inappropriate action differently, i.e. I allow for $x_{DM} \neq x_A$ and/or $y_{DM} \neq y_A$. The advisor is *unbiased* if $x_{DM} = x_A$ and $y_{DM} = y_A$. If player i cares relatively more about the loss from falsely accepting compared to the loss from falsely rejecting the hypothesis than player j, i.e. if $\frac{y_i}{x_i} > \frac{y_j}{x_j}$, then player i is *more reluctant* to switch than player j. Conversely, if player i cares relatively more about the loss from falsely rejecting compared to the loss from falsely accepting than player j, i.e. if $\frac{y_i}{x_i} < \frac{y_j}{x_j}$, then player i is *more enthusiastic* about switching than player j.

Each player minimises his expected loss. In addition, the advisor also takes into account the expected cost of testing.

I study the case in which the DM's preferences are such that her default choice is a_R, i.e. at the prior belief q it holds that

$$\begin{aligned} E(L_{DM}(a_R, \omega)) &< E(L_{DM}(a_S, \omega)) \\ &\Leftrightarrow \\ qx_{DM} &< (1-q)\, y_{DM} \end{aligned}$$

The case in which preferences are such that the DM prefers a_S at the prior belief would generate a "mirror image" of the case studied here.

Information Acquisition - The advisor has access to the following noisy but informative testing technology which generates a binary outcome $o \in \{(+),(-)\}$:

$$Pr\left(o=(+)\,|\omega=\omega^T\right)\equiv p_1$$

$$Pr\left(o=(-)\,|\omega=\omega^F\right)\equiv p_2,$$

where $\frac{1}{2}<p_i<1$ for $i=1,2$. The realisation of test outcomes is independent conditional on the state of the world. There is no further restriction on the relative size of p_2 and p_1, i.e. the accuracy of the test may depend on the state. The advisor incurs a private cost $C\geq 0$ to conduct one such test. It will be convenient to express the advisor's total loss relative to x_A, i.e. relative to the loss from choosing a_R if the hypothesis is true. In what follows, I will use $c\equiv\frac{C}{x_A}$ to capture the cost of testing relative to this loss x_A.

Structure of the Game - I consider two different settings. Either testing is observable to the DM or testing is hidden and the advisor communicates selected findings to the DM. In either case, the testing technology, cost of testing, each player's preferences and the structure and timing of the game are common knowledge. In addition, in either case the DM can neither commit ex ante to how she will use the additional information nor can she commit to paying any monetary rewards based on how much or on what she learns.[3]

2.2.2 Observable Testing

Timing and Information - Initially, Nature draws a state of the world, either ω^T or ω^F, which is not observed by either player. In each period $n=1,2$, the advisor decides whether or not to test. If he tests Nature draws an outcome, i.e. $o_n=(+)$ or $o_n=(-)$. Otherwise, Nature does not draw an outcome, i.e. $o_n=(\emptyset)$. Nature's draw is observed by both players. At the end of period 2, the DM takes an action and payoffs are realised.

Histories - Let h_n denote the *history of test outcomes* at the end of period $n=1,2$. In particular, $h_2=(o_1o_2)$ and $h_1=(o_1)$ where $o_n\in\{(+),(-),(\emptyset)\}$. Note that the history of test outcomes h_n characterises the complete history of play up to period n except for Nature's draw of the state of the world.

Strategies - Let the advisor's testing strategy be characterised by $T_A=\left(t_1,t_2^+,t_2^-,t_2^\emptyset\right)$, where $t_1\in\{Y,N\}$ denotes the advisor's decision to test (Y) or not to test (N) in period 1 and $t_2^{h_1}\in\{Y,N\}$ denotes

[3]This captures situations in which no contracts can be written and no credible promises regarding future actions can be made.

the decision in period 2 conditional on the history of test outcomes at the end of period 1, i.e. conditional on $h_1 \in \{+, -\emptyset\}$. Let the DM's decision rule be characterised by Z_{DM}, where Z_{DM} is a set of histories of outcomes at the end of period 2 such that the DM chooses a_S if and only if $h_2 \in Z_{DM}$.

Equilibrium Concept - The solution concept is a Perfect Bayesian Nash equilibrium in pure strategies. An equilibrium is a vector (Z_{DM}, T_A) such that Z_{DM} is a best response to T_A and T_A is a best response to Z_{DM}.

2.2.3 Hidden Testing

Timing and Information - Initially, Nature draws a state of the world, either ω^T or ω^F, which is not observed by either player. In period $n = 1,2$, the advisor decides whether or not to test. If he tests then Nature draws an outcome, i.e. $o_n = (+)$ or $o_n = (-)$. Otherwise, Nature does not draw an outcome, i.e. $o_n = (\emptyset)$. The advisor observes Nature's draw, but the DM observes neither Nature's draw nor the advisor's move. At the end of period 2, the advisor sends a message m to the DM and then the DM takes an action and payoffs are realised.

Histories - Let h_n denote the *history of test outcomes* at the end of period n, just as in the case of observable testing.

Communication - The message space is given by $m \in \{(\emptyset), (-), (+), (-+), (++), (--)\}$. The advisor can only communicate that he has tested by revealing an outcome and if he reveals an outcome he cannot communicate when this outcome was discovered. For any given history of test outcomes h_2, the advisor is further restricted regarding the messages he can send because he can only include outcomes he has observed, but he cannot invent outcomes. However, he is able to omit outcomes. For example, if $h_2 = (-+)$ then the set of feasible messages is $m \in \{(\emptyset), (-), (+), (-+)\}$. This assumption reflects that when checking scientific work it is very difficult to know whether the scientist conducted additional experiments, but it is relatively easy to discover if he faked the data.

I do not allow the advisor to report at the end of the first period. Thus, he cannot credibly convey that he has conducted fewer than two tests by the timing of his report. This captures that the timing of the report may not be verifiable by the DM.

Strategies - Let the advisor's strategy be denoted by (T_A, M_A), where $T_A = \left(t_1, t_2^+, t_2^-, t_2^\emptyset\right)$ characterises his testing strategy, as in the case of observable testing, and M_A denotes his reporting strategy. Let M_A map any history of test outcomes h_2 to a message m subject to the constraints mentioned above.

Unlike in the case of observable testing, the DM cannot condition her action on the history of test outcomes h_2. Instead, she can only condition her choice on the message received. The DM's message-based decision rule is characterised by Σ_{DM}, where Σ_{DM} is a set of messages such that the DM chooses a_S if and only if $m \in \Sigma_{DM}$. The DM has a belief assessment b which for any message m gives a probability distribution over histories of outcomes h_2 conditional on each state of the world, i.e. to each history of test outcomes h_2 she assigns beliefs $Pr_{\omega^T}(h_2|m)$ conditional on the state being $\omega = \omega^T$ and $Pr_{\omega^F}(h_2|m)$ conditional on the state being $\omega = \omega^F$ for any given message m.

Equilibrium Concept - The solution concept is a Perfect Bayesian Nash equilibrium in pure strategies. An equilibrium is a vector $(\Sigma_{DM}, T_A, M_A, b)$ such that Σ_{DM} is optimal given b, (T_A, M_A) is a best response given Σ_{DM} and belief assessment b is consistent.[4]

2.3 No Conflict of Interest - Single Player Benchmark

In this benchmark, I analyse the situation in which a single player chooses whether or not to test and which action to take. I characterise how the player behaves on the equilibrium path contingent on all possible draws of test outcomes. His behaviour depends on his preferences and on the cost of testing, for a given prior belief that the hypothesis is true. This benchmark serves as a useful reference point for future analysis because the single player's behaviour matches the behaviour of the DM and the advisor when their preferences are aligned, irrespective of whether testing is observable or hidden.

The single player's preferences are captured by $\frac{y}{x}$. I allow for the full range of preferences $\frac{y}{x} \in [0, \infty)$, i.e. the default action may be either a_R or a_S. I focus on the player's behaviour in equilibrium rather than his strategy because I am interested in comparing expected payoffs achieved in different settings.[5] To characterise the player's behaviour, I need to define whether or not he tests in the first period, whether or not he continues testing in the second period conditional on the outcome of the first test and which action he chooses given the outcomes. It is useful to introduce the following notation. $T0$, $T1$ and $T2$ refer to the number of tests which the player conducts at most. A superscript on $T2$ denotes for which outcome of the first test the player continues testing. A subscript denotes the set of test outcomes for which the player chooses a_S. Equilibrium behaviour of the player is thus described as follows:

[4] I conjecture that the payoffs derived in the game of hidden testing are unique if the players are sequentially rational and if the advisor is truth-leaning, i.e. if the advisor reveals an outcome whenever he is indifferent between revealing and hiding.

[5] The derivation of the player's optimal strategy can be found in the Appendix.

$T0$ The player does not test. He chooses his default action.

$T1_{(+)}$ The player tests in the first period, but never in the second. He chooses a_S if and only if the test outcome is positive.

$T2^{+}_{(++)}$ The player tests in the first period. If the outcome is negative he chooses a_R. If the outcome is positive, he tests again. If the second outcome is also positive he chooses a_S. Otherwise he chooses a_R.

$T2^{-}_{(+),(-+)}$ The player tests in the first period. If the outcome is positive he chooses a_S. If the outcome is negative he tests again. If the second outcome is positive he chooses a_S, but if the second outcome is also negative he chooses a_R.

To characterise how the player's behaviour in equilibrium depends on his preferences and on the cost of testing, it is useful to first define threshold levels of the cost of testing c.

Definition 1

- *At* $c = \overline{c}^{S}_{01}\left(\frac{y}{x}\right)$*, the player is indifferent between* $T0$ *and* $T1_{(+)}$ *when his default action is* a_S

$$\overline{c}^{S}_{01}\left(\frac{y}{x}\right) = (1-q)\,p_2\frac{y}{x} - (1-p_1)\,q$$

- *At* $c = \overline{c}^{S}_{02}\left(\frac{y}{x}\right)$*, the player is indifferent between* $T0$ *and* $T2^{-}_{(+),(-+)}$ *when his default action is* a_S

$$\overline{c}^{S}_{02}\left(\frac{y}{x}\right) = \frac{p_2\,(1-q)\,\frac{y}{x} - (1-p_1)^2\,q}{1+p_2\,(1-q)+(1-p_1)\,q}$$

- *At* $c = \overline{c}^{-}_{12}\left(\frac{y}{x}\right)$ *the player is indifferent between* $T1_{(+)}$ *and* $T2^{-}_{(+),(-+)}$

$$\overline{c}^{-}_{12}\left(\frac{y}{x}\right) = \frac{p_1\,(1-p_1)\,q - p_2\,(1-p_2)\,(1-q)\,\frac{y}{x}}{p_2\,(1-q)+(1-p_1)\,q}$$

- *At* $c = \overline{c}^{R}_{01}\left(\frac{y}{x}\right)$*, the player is indifferent between* $T0$ *and* $T1_{(+)}$ *when his default action is* a_R

$$\overline{c}^{R}_{01}\left(\frac{y}{x}\right) = qp_1 - (1-q)\,(1-p_2)\,\frac{y}{x}$$

- *At* $c = \overline{c}^{R}_{02}\left(\frac{y}{x}\right)$*, the player is indifferent between* $T0$ *and* $T2^{+}_{(++)}$ *when his default action is* a_R

$$\overline{c}^{R}_{02}\left(\frac{y}{x}\right) = \frac{p_1 q - (1-p_2)^2 (1-q)\frac{y}{x}}{1+(1-p_2)(1-q)+p_1 q}$$

- *At* $c = \overline{c}^{+}_{12}\left(\frac{y}{x}\right)$ *the player is indifferent between* $T1_{(+)}$ *and* $T2^{+}_{(++)}$

$$\overline{c}^{+}_{12}\left(\frac{y}{x}\right) = \frac{p_2 (1-p_2)(1-q)\frac{y}{x} - p_1 (1-p_1) q}{(1-p_2)(1-q)+p_1 q}$$

Lemma 1 [Behaviour in Equilibrium - Single Player]

The player's behaviour in equilibrium is described by

- $T2^{-}_{(+),(-+)}$ *if and only if* $c \leq min\left\{\overline{c}^{-}_{12}\left(\frac{y}{x}\right), \overline{c}^{S}_{02}\left(\frac{y}{x}\right)\right\}$,
- $T1_{(+)}$ *if and only if* $c > max\left\{\overline{c}^{-}_{12}\left(\frac{y}{x}\right), \overline{c}^{+}_{12}\left(\frac{y}{x}\right)\right\}$ *and* $c \leq min\left\{\overline{c}^{S}_{01}\left(\frac{y}{x}\right), \overline{c}^{R}_{01}\left(\frac{y}{x}\right)\right\}$,
- $T2^{+}_{(++)}$ *if and only if* $c \leq min\left\{\overline{c}^{+}_{12}\left(\frac{y}{x}\right), \overline{c}^{R}_{02}\left(\frac{y}{x}\right)\right\}$,
- *and* $T0$ *otherwise.*

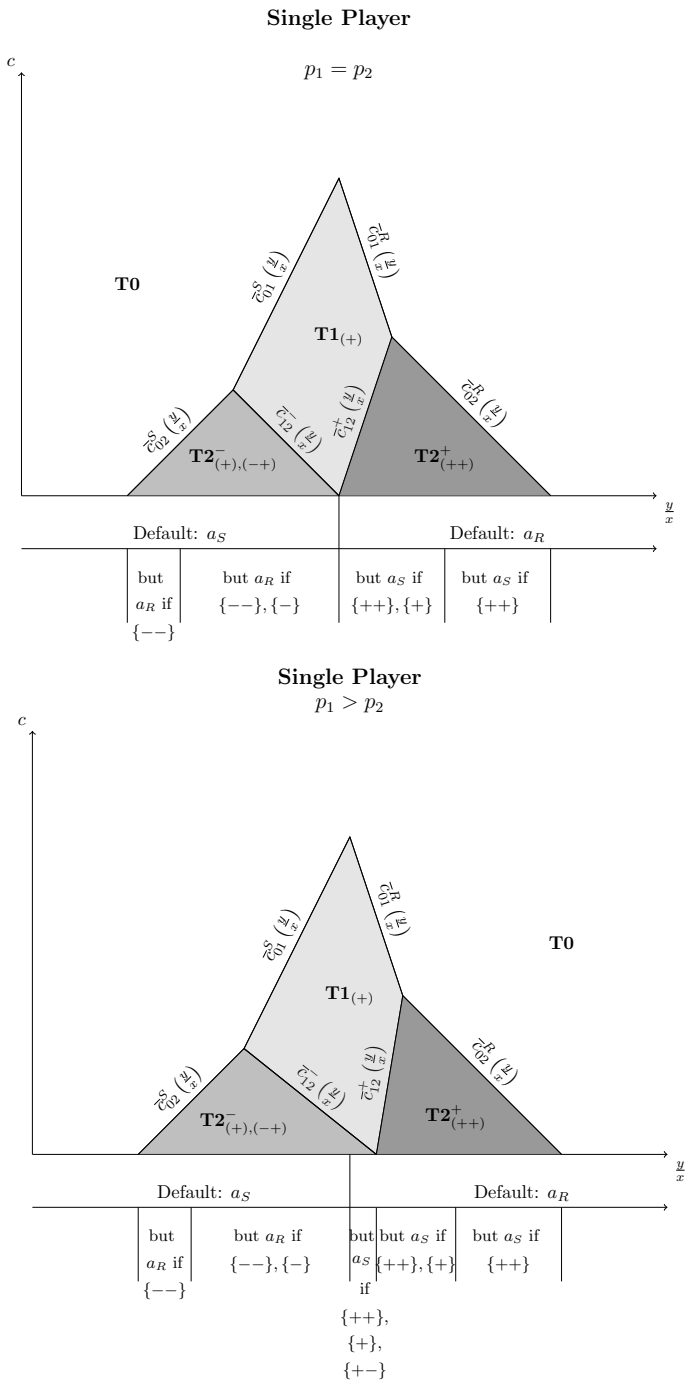

Figure 2.1: These figures illustrate the single player's behaviour in equilibrium for combinations of cost c and preferences as characterised by $\frac{y}{x}$, holding fixed the prior belief q and the testing technology as described by p_2 and p_1. Moving along the horizontal axis, the player becomes more reluctant. The horizontal axis is divided into ranges for which a different decision rule is optimal. The optimal decision rule as a function of the number of positive and negative outcomes is described underneath.

How the player's behaviour in equilibrium depends on the cost of testing and on his preferences is illustrated in Figure 2.1. Consider the situation in which the accuracy of the test is independent of the state of the world, i.e. $p_1 = p_2$ (upper panel). Suppose the player's loss x from choosing a_R if the hypothesis is true is fixed. As the player's loss y from choosing a_S if the hypothesis is false increases, he becomes more reluctant to switch. For sufficiently high y, his default action is a_R. In what follows, I will focus on the range of preferences such that the default action is a_R, but the reasoning applies to the "mirror image" if the default action is a_S.

The higher the loss y from choosing a_S if the hypothesis is false, the lower is the player's willingness to test. This is because testing is only worthwhile if positive evidence can lead him to choose a_S, but such evidence is not very valuable when choosing a_S is associated with the risk of a high loss in case the hypothesis is false. Suppose the player is relatively enthusiastic about switching (low y) such that he would choose a_S if and only if he were to observe either two positive outcomes or one positive outcome in a single test. Such a player benefits from testing in the first period because one positive outcome will convince him that a_S yields a lower loss in expectation. If the cost is low, he will prefer to check whether a first positive outcome is robust to an additional test and will choose a_S only if he observes two positive outcomes. The higher the loss y from choosing a_S if the hypothesis is false, the more cautious the player is about choosing a_S after one positive outcome in a single test and thus the more willing he is to conduct an additional test. Now suppose that the player is more reluctant (higher y) such that he would choose a_S if and only if he were to observe two positive outcomes, but not yet if he were to observe one positive outcome in a single test. In this case, he will only start testing if he is willing to incur the cost of a second test following a positive outcome. He will never stop testing after a positive outcome because one positive outcome can never convince him to choose a_S. Finally, suppose that the player is very reluctant (high y). Then he will never test because there is no test outcome that could convince him to choose a_S.

Next, suppose that the test accuracy is state-dependent and a false negative test outcome is more likely than a false positive test outcome, i.e. $p_1 > p_2$ (lower panel). This implies that one positive and one negative outcome are evidence in favour of the hypothesis. What is special about this situation is that there is a range of preferences such that the player chooses a_R as a default choice, yet he is rather enthusiastic about switching and would choose a_S if he were to observe one positive and one negative outcome. Such a player will never conduct an additional test following a positive test outcome because he would choose a_S irrespective of what outcome an additional test yields. However, if the first test

is negative he may check whether a second test yields a positive outcome. A positive outcome in the second test is stronger evidence in favour than the negative outcome in the first test is evidence against the hypothesis. He will only conduct a second test if the cost is sufficiently low. For a higher cost he will either conduct one test or no test at all.

2.4 Conflict of Interest with Observable Testing

This section analyses the interaction between a DM and an advisor when the DM relies on the advisor to conduct tests and test outcomes are observable to both the DM and the advisor. A conflict of interest arises for two reasons: i) only the advisor incurs the cost of testing and ii) the advisor may have different preferences than the DM. I derive the behaviour of the advisor and the DM contingent on all possible draws of test outcomes in the pure-strategy Nash equilibrium of this game.

In equilibrium, the DM does not choose her decision rule strategically. She observes test outcomes and acts optimally given this evidence, independent of the advisor's testing strategy. However, the advisor chooses his testing strategy as a best response to how the DM will act based on the evidence observed.

I restrict attention to the range of DM's preferences such that she chooses a_R as a default. To simplify the exposition, it is helpful to divide the range of DM's preferences into sections, such that the DM's optimal decision rule is constant within each section but differs across sections. I will then assign labels to each section to describe the decision rule the DM will follow. In particular, I will label the DM either progressive, moderate or conservative where the conservative DM is the one who is willing to switch for the smallest set of possible test outcomes, as illustrated in Figure 2.2. In what follows, I will ignore types of DM who would not switch for any evidence. I will adopt the notation introduced in Section 2.3 to characterise players' behaviour in equilibrium.

Definition 2a

- *A **progressive** DM follows decision rule*

$$Z_{DM} = \{(+-), (-+), (\emptyset+), (+\emptyset), (++)\}$$

and her preferences satisfy

$$\frac{q}{1-q} \leq \frac{y_{DM}}{x_{DM}} < \frac{qp_1(1-p_1)}{(1-q)p_2(1-p_2)}.$$

Note that $p_1 > p_2$ is a necessary condition for the DM to be progressive.

- *A **moderate** DM follows decision rule*

$$Z_{DM} = \{(\emptyset+), (+\emptyset), (++)\}$$

and her preferences satisfy

$$max\left\{\frac{q}{1-q}, \frac{qp_1(1-p_1)}{(1-q)p_2(1-p_2)}\right\} \leq \frac{y_{DM}}{x_{DM}} < \frac{qp_1}{(1-q)(1-p_2)}.$$

- *A **conservative** DM follows decision rule*

$$Z_{DM} = \{(++)\}$$

and her preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)} \leq \frac{y_{DM}}{x_{DM}} < \frac{qp_1}{(1-q)(1-p_2)^2}.$$

Progressive (requires $p_1 > p_2$)						Moderate						Conservative					
	$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$
a_S	+	+	a_S	+	+	a_S	+	+	a_S	+	+	a_S	+	+	a_R	+	+
a_S	+	−		+	−	a_R	+	−		+	−	a_R	+	−		+	−
a_S	−	+	a_R	−	+	a_R	−	+	a_R	−	+	a_R	−	+	a_R	−	+
a_R	−	−		−	−	a_R	−	−		−	−	a_R	−	−		−	−

Figure 2.2: The tables illustrate the optimal decision rule for each type of DM. The first half of each table illustrates the optimal action if the outcome of each test is known. The second half illustrates the optimal action if only the outcome of the first test is known.

Lemma 2 [Behaviour in Equilibrium - Observable Testing]

Suppose test outcomes are observable.

Suppose $p_1 > p_2$ *and the DM is* ***progressive.*** *Then*

- *the equilibrium with* $T2^{-}_{(+),(-+)}$ *exists if and only if* $c \leq \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$,
- *the equilibrium with* $T1_{(+)}$ *exists for* $\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right) < c \leq \overline{c}^{R}_{01}\left(\frac{y_A}{x_A}\right)$,
- *otherwise only the equilibrium with* $T0$ *exists.*

Suppose the DM is ***moderate.*** Then

- *the equilibrium with* $T2^{+}_{(++)}$ *exists if and only if* $c \leq min\left\{\overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)\right\}$,
- *the equilibrium with* $T1_{(+)}$ *exists if and only if* $\overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) \leq c < \overline{c}^{R}_{01}\left(\frac{y_A}{x_A}\right)$,
- *otherwise only the equilibrium with* $T0$ *exists.*

Suppose the DM is ***conservative.*** Then

- *the equilibrium with* $T2^{+}_{(++)}$ *exists if and only if* $c \leq \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$,
- *otherwise only the equilibrium with* $T0$ *exists.*

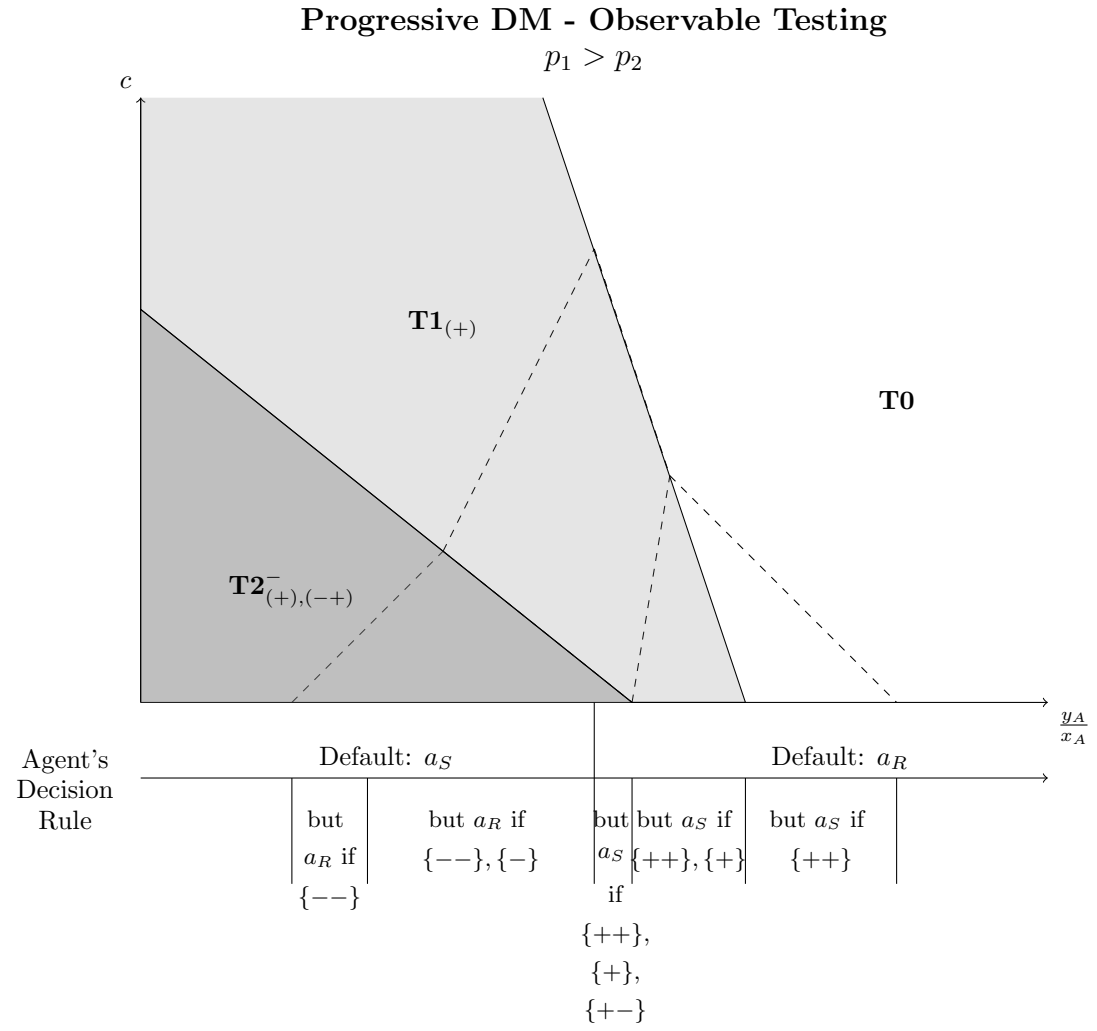

Figure 2.3: This figure shows behaviour in the pure-strategy Nash equilibrium when the DM is progressive (see Def 2a) for combinations of cost c and the advisor's preferences as captured by $\frac{y_A}{x_A}$. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

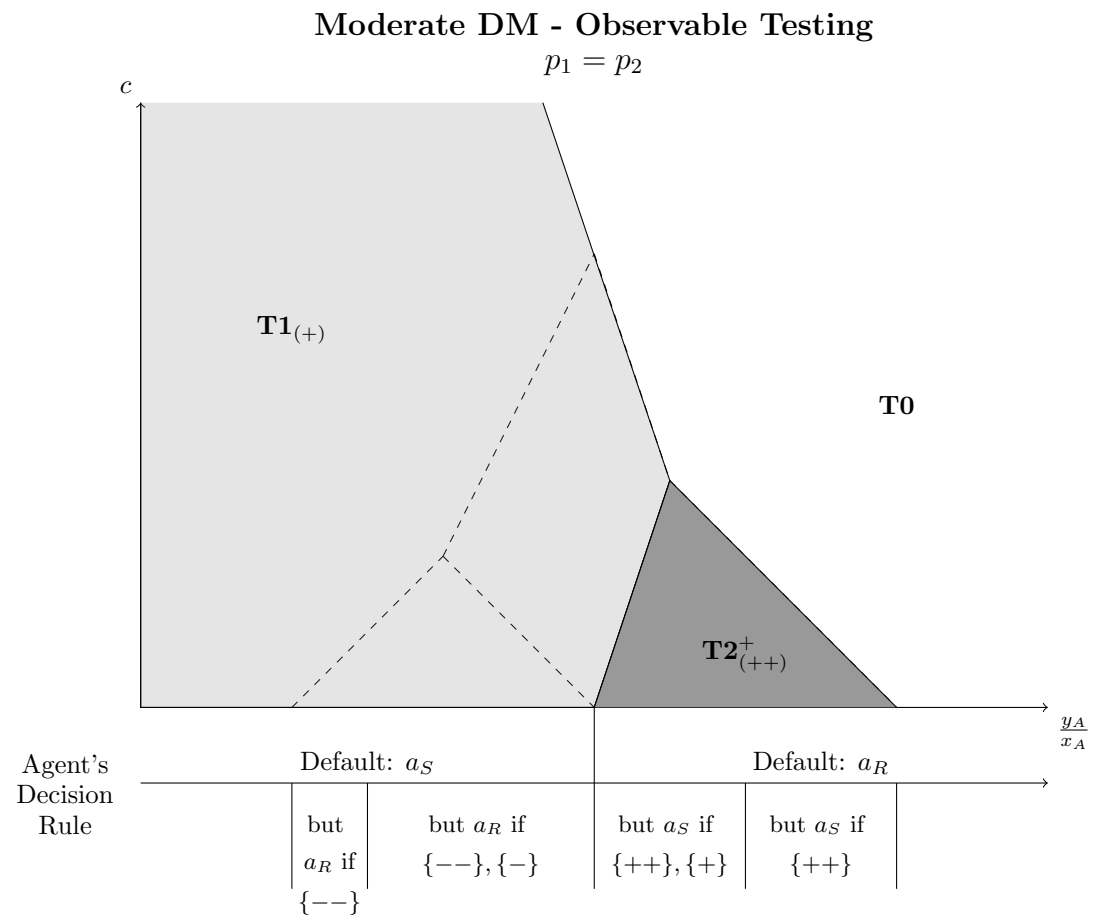

Figure 2.4: This figure shows behaviour in the pure-strategy Nash equilibrium when the DM is moderate (see Def 2a) for combinations of cost c and the advisor's preferences as captured by $\frac{y_A}{x_A}$. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

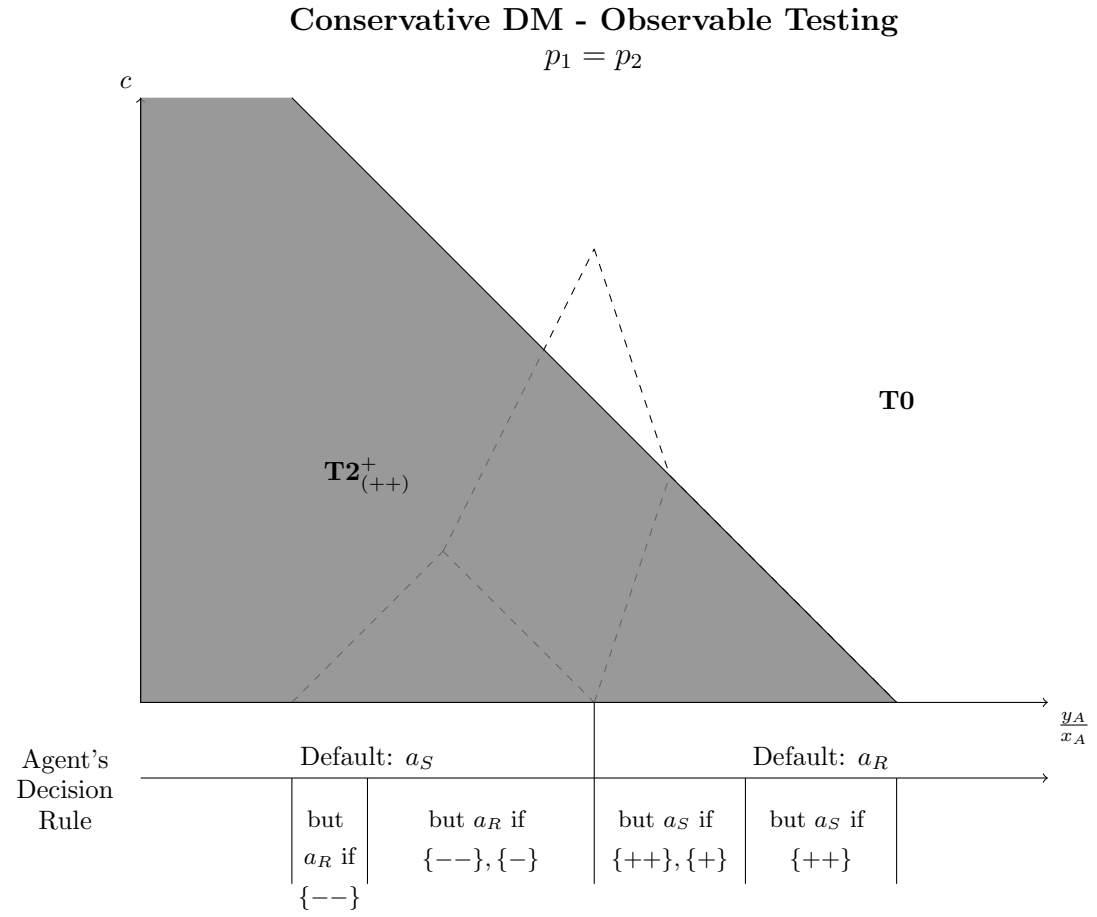

Figure 2.5: This figure shows behaviour in the pure-strategy Nash equilibrium when the DM is conservative (see Def 2a) for combinations of cost c and the advisor's preferences as captured by $\frac{y_A}{x_A}$. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

To describe what happens in equilibrium, I will fix the type of DM the advisor interacts with and explain the strategic behaviour of an advisor who is more reluctant than this DM and the behaviour of an advisor who is more enthusiastic than this DM. If the advisor is unbiased then he will choose his testing strategy as if he could also choose the action himself and the behaviour of both players in equilibrium is as in the case of a single player (see Section 2.3).

Suppose the test has the feature that a false negative outcome is more likely than a false positive outcome, i.e. $p_1 > p_2$, which implies that observing one positive and one negative outcome is evidence in favour of the hypothesis being true. Consider a **progressive** DM, a DM who would switch based on one positive and one negative outcome or stronger evidence in favour of the hypothesis, as illustrated in Figure 2.3. Suppose the advisor is more enthusiastic about switching than the progressive DM, which means that the advisor would choose a_S as the default if he were to decide. This advisor will conduct tests such as to find evidence in favour of the hypothesis. If the evidence is sufficient for the DM to switch, then the more enthusiastic advisor also approves of the decision to switch. The advisor will test at least once provided the cost is sufficiently low, because one positive outcome is sufficient evidence in favour of the hypothesis for the DM to switch. For a very low cost, the advisor may even benefit from testing again if the first test is negative, because if the second test were positive then the DM would also be convinced to switch. Furthermore, suppose that the advisor is more reluctant than the DM. He also seeks evidence in favour of the hypothesis, but the evidence may be such that the DM is convinced to switch when the advisor himself is not yet convinced. Suppose the advisor is only slightly more reluctant such that if he were to choose the action he would not switch based on one positive and one negative outcome, but he would switch for one positive outcome in a single test or two positive outcomes in two tests. This advisor will test at most once. If the outcome is negative, the advisor does not want the DM to switch irrespective of what a second test would yield and without a further test the DM will choose a_R. If the outcome is positive then the DM will switch, which is in line with what the advisor would do. While the advisor may be willing to incur the cost to see if this positive outcome is confirmed by a further test, he knows that the DM will switch regardless of the outcome of the second test and therefore he will never conduct a second test. Finally, suppose that the advisor is much more reluctant than the DM such that he would want to switch if and only if he were to observe two positive outcomes. Then the advisor will never test. This is because he can never achieve that the DM will only switch if two outcomes are positive. If the first test outcome is positive, the DM will switch even if the second test contradicts the first. If the first

test outcome is negative then the DM will choose a_R, but this is also the DM's default choice if he does not test.

Consider a **moderate** DM, a DM who is willing to switch based on one positive outcome in a single test or two positive outcomes in two tests, as illustrated in Figure 2.4. Suppose that the advisor is more enthusiastic than the DM. Then the advisor will test at most once. If the test outcome is positive, then the DM will choose a_S, which is in the advisor's interest. If the test outcome is negative, then the DM can never be swayed to choose a_S whatever the second test outcome may yield. Now suppose the advisor is more reluctant than the DM and only wants to switch if and only if he were to observe two positive test outcomes. Then the advisor will start testing only if he is willing to incur the cost for a second test following a positive outcome. He will never stop testing after a positive outcome in the first period. This is because the DM will choose a_S based on one positive outcome, which is not in the advisor's interest. But if the advisor conducts an additional test after a positive outcome the DM acts in the advisor's interest and only switches if both outcomes are positive.

Lastly, consider a **conservative** DM, a DM who is willing to switch if and only if two positive outcomes are observed, as illustrated in Figure 2.5. If the advisor is more enthusiastic than the DM, then the advisor will start testing only if he is willing to incur the cost for a second test following a positive outcome. He will never stop testing after a positive outcome. That is because one test is insufficient to sway the DM to switch. If the advisor is more reluctant than the DM he will never test because there is no outcome such that he himself would want to switch and by not testing he can guarantee that the DM chooses a_R.

2.5 Conflict of Interest with Hidden Testing

This section analyses the interaction between a DM and an advisor when the DM can neither observe test outcomes nor the advisor's decision to test. Instead, she has to base her choice on the advisor's report in which the advisor can selectively disclose outcomes. I will first characterise players' behaviour contingent on all possible draws of test outcomes in a pure-strategy Bayesian Nash equilibrium of this game. Then I will identify the circumstances in which the DM is better off when testing is hidden rather than observable.

In equilibrium, the advisor chooses both a testing strategy and a reporting strategy. The DM bases

her choice on the advisor's report and holds a set of beliefs about what the advisor may have observed conditional on the report she has received. There are multiple equilibria. I conjecture that the equilibrium selected in Lemma 3 is the one which maximises the Advisor's expected payoff. The remainder of my analysis is based on this selection of equilibria.

Under hidden testing, the enlarged scope for the advisor to influence the DM's decision requires further partitioning of the range of preferences for which the DM is considered moderate. In particular, consider the following behaviour of the advisor: He tests once and if the outcome is positive then he stops testing and reports a single positive outcome. If the outcome is negative then he tests again. If both test outcomes are negative then he reports two negative outcomes, but if the second test outcome is positive then he reports a single positive outcome, i.e. he hides one negative outcome. A DM who anticipates this behaviour forms the belief that the report of a single positive outcome means that the advisor observed one positive outcome in a single test or that he observed a negative outcome in the first test and a positive outcome in the second test. Whether or not a moderate DM will choose a_S based on the report of a single positive outcome is not clear. Therefore, I partition the range of preference of a moderate DM such that a moderate DM who chooses a_R is labelled moderate leaning towards conservative and a moderate DM who chooses a_S is labelled moderate leaning towards progressive (see Figure 2.6):

Definition 2b

- *A **DM is moderate leaning towards progressive** if and only if*

$$max\left\{\frac{q}{1-q}, \frac{qp_1(1-p_1)}{(1-q)p_2(1-p_2)}\right\} \leq \frac{y_{DM}}{x_{DM}} < \frac{q(p_1+p_1(1-p_1))}{(1-q)(1-p_2+p_2(1-p_2))}.$$

- *A **DM is moderate leaning towards conservative** if and only if*

$$\frac{q(p_1+p_1(1-p_1))}{(1-q)(1-p_2+p_2(1-p_2))} \leq \frac{y_{DM}}{x_{DM}} < \frac{qp_1}{(1-q)(1-p_2)}.$$

Lemma 3 [Behaviour in Equilibrium - Hidden Testing]

Suppose test outcomes are hidden from the DM and the advisor can selectively disclose outcomes.

Suppose the DM is ***progressive*** *(and $p_1 > p_2$)* ***or moderate leaning towards progressive.*** *Then*

- *an equilibrium with $T2^{-}_{(+),(-+)}$ exists if $c < \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$,*
- *an equilibrium with $T1_{(+)}$ exists if $max\left\{\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right)\right\} < c < \overline{c}^{R}_{01}\left(\frac{y_A}{x_A}\right)$,*
- *an equilibrium with $T2^{+}_{(++)}$ exists if $c < min\left\{\overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)\right\}$,*
- *otherwise an equilibrium with $T0$ exists.*

Suppose the DM is ***moderate leaning towards conservative.*** *Then*

- *an equilibrium with $T1_{(+)}$ exists if $max\left\{\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right)\right\} < c < \overline{c}^{R}_{01}\left(\frac{y_A}{x_A}\right)$,*
- *an equilibrium with $T2^{+}_{(++)}$ exists if either $c < min\left\{\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)\right\}$ or $c < min\left\{\overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)\right\}$,*
- *otherwise an equilibrium with $T0$ exists.*

Suppose the DM is ***conservative.*** *Then*

- *an equilibrium with $T2^{+}_{(++)}$ exists if $c < \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$,*
- *otherwise an equilibrium with $T0$ exists.*

Moderate leaning towards progressive									Moderate leaning towards reluctant								
	$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$		$t=1$	$t=2$
a_S	+	+	a_S	+	+	a_S	+	+	a_S	+	+	a_S	+	+	a_R	+	+
a_R	+	−		+	−		+	−	a_R	+	−		+	−		+	−
a_R	−	+	a_R	−	+		−	+	a_R	−	+	a_R	−	+		−	+
a_R	−	−		−	−	a_R	−	−	a_R	−	−		−	−	a_R	−	−

Figure 2.6: The tables illustrate the optimal decision rule for a moderate leaning towards progressive and a moderate leaning towards conservative DM. The first section of each table illustrates the optimal action if the outcome of each test is known, the second section illustrates the optimal action if only the outcome of the first test is known, the final section illustrates the optimal action if it is known that either the first test was positive or the first test was positive and the second test was negative.

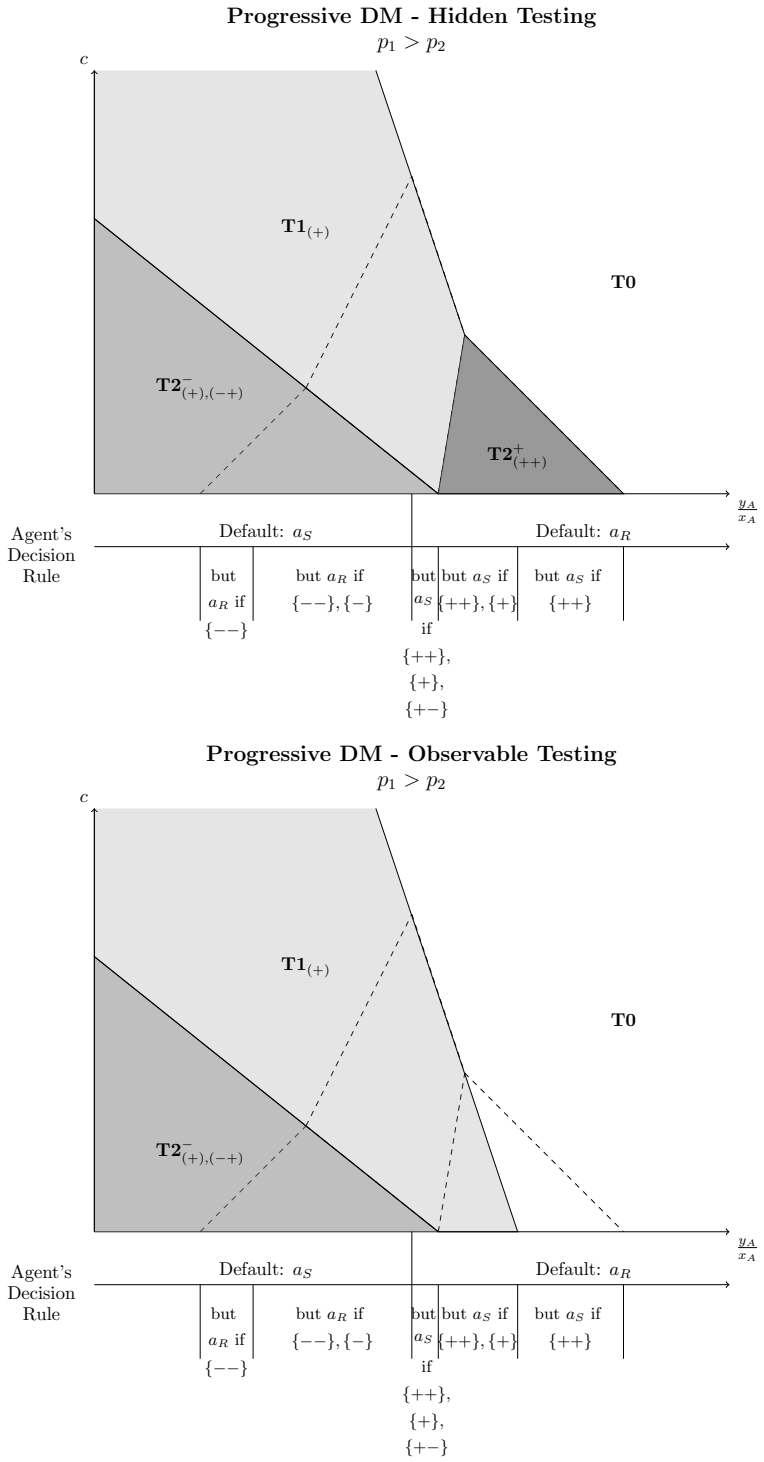

Figure 2.7: These figures illustrate behaviour in equilibrium given the DM is progressive for combinations of cost c and advisor's preferences as characterised by $\frac{y_A}{x_A}$, holding fixed the prior belief q and the testing technology as described by p_2 and p_1. The upper panel is drawn assuming testing is hidden and the advisor can selectively disclose outcomes, the lower panel is drawn assuming testing is observable. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

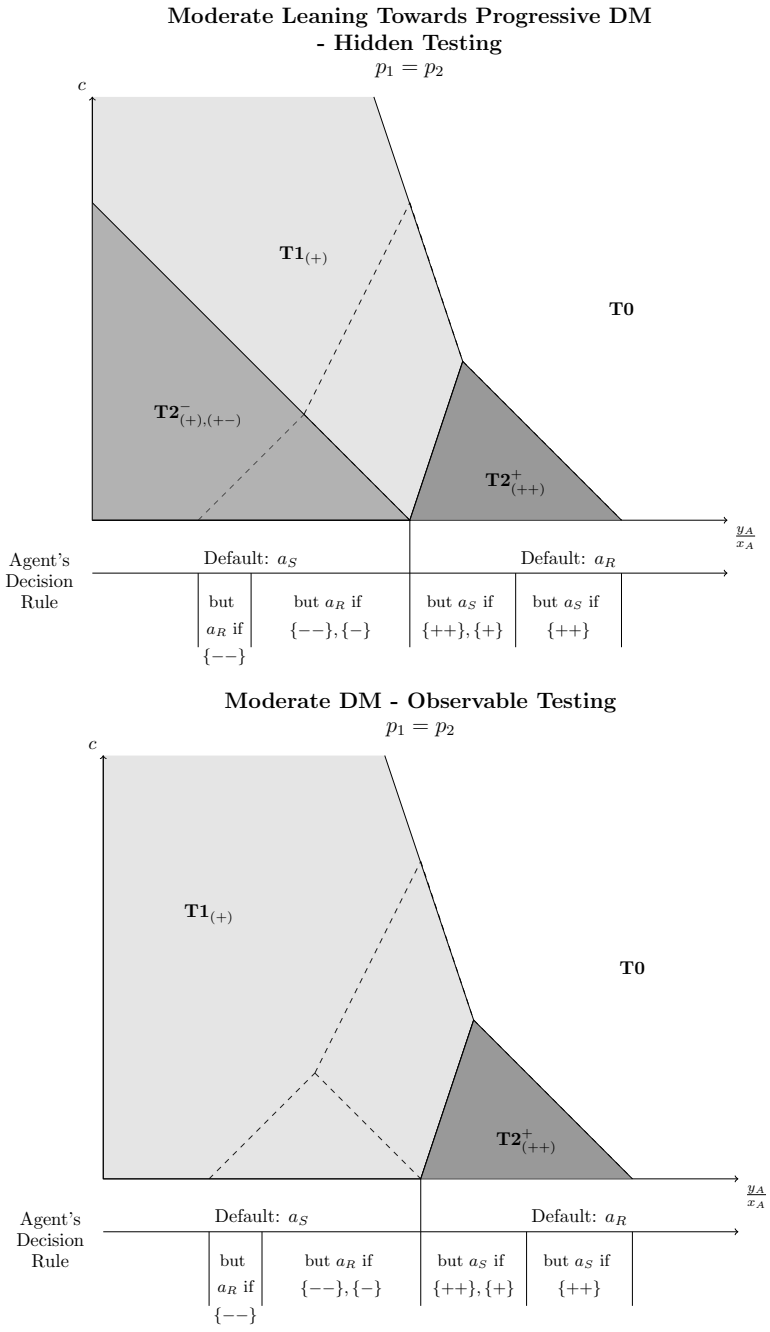

Figure 2.8: These figures illustrate behaviour in equilibrium given the DM is moderate leaning towards progressive for combinations of cost c and advisor's preferences as characterised by $\frac{y_A}{x_A}$, holding fixed the prior belief q and the testing technology as described by p_2 and p_1. The upper panel is drawn assuming testing is hidden and the advisor can selectively disclose outcomes, the lower panel is drawn assuming testing is observable. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

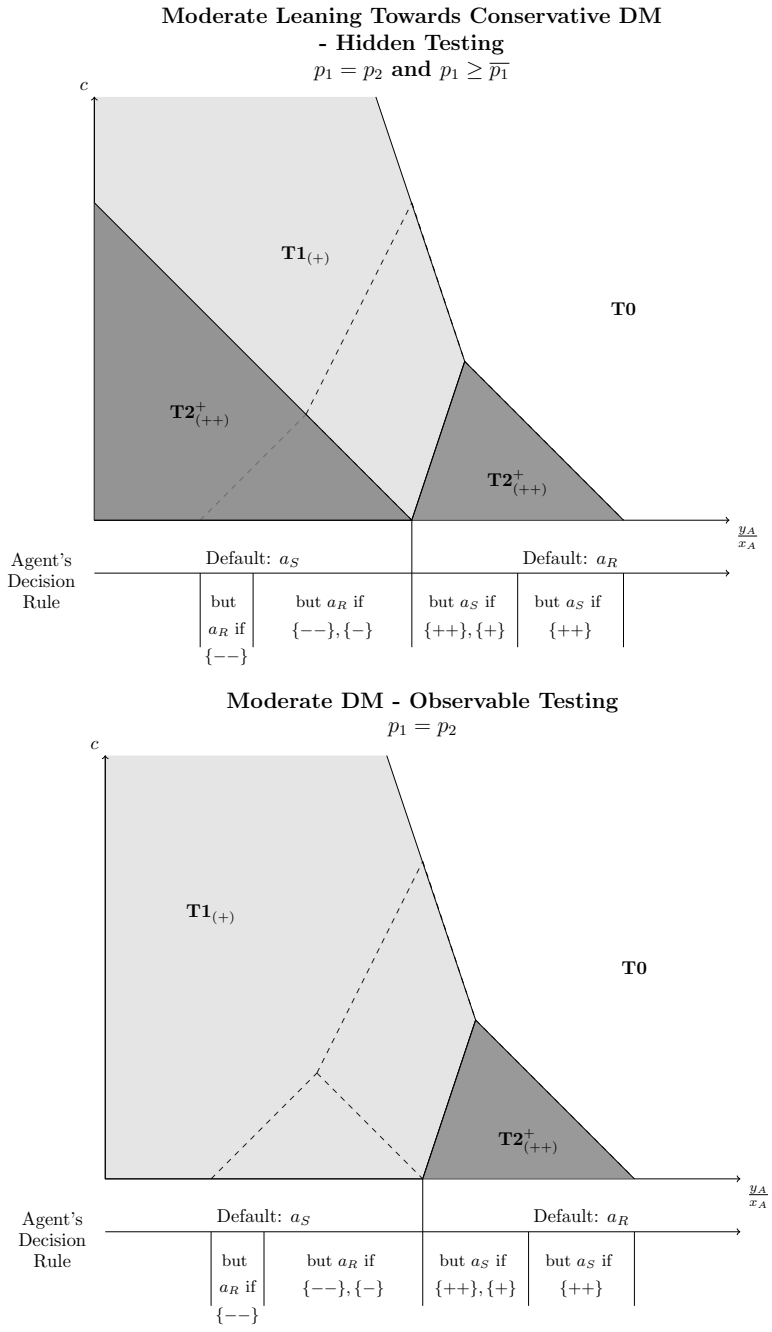

Figure 2.9: These figures illustrate behaviour in equilibrium given the DM is moderate leaning towards conservative for combinations of cost c and advisor's preferences as characterised by $\frac{y_A}{x_A}$, holding fixed the prior belief q and the testing technology as described by p_2 and p_1. I assume $p_1 \geq \overline{p_1}$ where $\overline{p_1}$ is characterised by $\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}=0\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}=0\right)$. The upper panel is drawn assuming testing is hidden and the advisor can selectively disclose outcomes, the lower panel is drawn assuming testing is observable. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

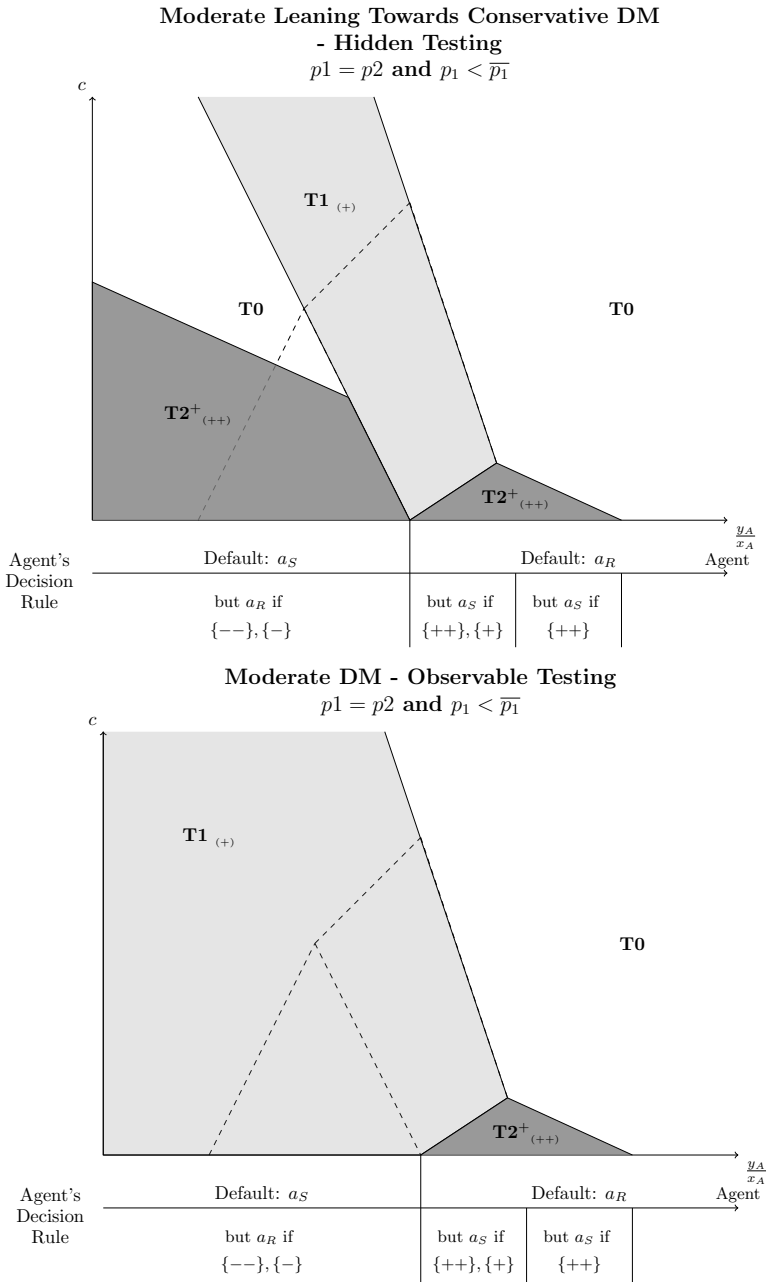

Figure 2.10: These figures illustrate behaviour in equilibrium given the DM is moderate leaning towards conservative for combinations of cost c and advisor's preferences as characterised by $\frac{y_A}{x_A}$, holding fixed the prior belief q and the testing technology as described by p_2 and p_1. I assume $p_1 < \overline{p_1}$ where $\overline{p_1}$ is characterised by $\overline{c}_{12}^{-}\left(\frac{y_A}{x_A} = 0\right) = \overline{c}_{02}^{R}\left(\frac{y_A}{x_A} = 0\right)$. The upper panel is drawn assuming testing is hidden and the advisor can selectively disclose outcomes, the lower panel is drawn assuming testing is observable. The range of the advisor's preferences (along the horizontal axis) is divided into sections for which a different decision rule is optimal for the advisor. His optimal decision rule as a function of the number of positive and negative outcomes is described underneath. The dotted lines indicate the relevant boundaries of regions for the case of a single player (see Section 2.3).

I will focus on situations in which players' behaviour differs from their behaviour under observable testing.

Suppose the test has the feature that a false negative outcome is more likely than a false positive outcome, i.e. $p_1 > p_2$. First, consider a ***progressive*** DM, as illustrated in Figure 2.7. Under hidden testing, an advisor who is more reluctant than the DM may be tempted to hide some positive evidence. Provided the cost is low, such an advisor will test in the first period and then again following a positive outcome. If the second outcome is positive then he is himself convinced that a_S is optimal and therefore he reveals two positive outcomes to the DM, inducing the DM to act in line with his interests. Otherwise, the advisor will report a single negative outcome, hiding a positive outcome if necessary. This leads the DM to stick with a_R, again in line with the advisor's interest. Although the DM suspects that the advisor may be hiding some positive evidence when she receives a report of a single negative outcome, she will stick with a_R.

Furthermore, consider the case of the ***moderate leaning towards progressive*** DM, as illustrated in Figure 2.8. An advisor who is more enthusiastic than the DM may benefit from hiding negative evidence. Provided the costs is sufficiently low, such an advisor will test in the first period and again in the second period following a negative outcome. If the first outcome is positive and the second is negative then he will reveal only a single positive outcome. Otherwise, he will reveal all outcomes. When the DM receives a report of a single positive outcome then she anticipates that the advisor either discovered a positive outcome on the first test or he discovered a negative outcome on the first test and a positive outcome on the second test. By definition of a moderate leaning towards progressive DM, the report of a single positive outcome still induces the DM to switch. If two outcomes are negative then it does not matter if the advisor reveals these outcomes or if he pretends not to have tested, in either case the DM will stick to a_R.

Next, consider the case of the ***moderate leaning towards conservative*** DM, as illustrated in Figures 2.9 and 2.10. If the cost is sufficiently low, an advisor who is more enthusiastic is tempted to test in the first period and report the outcome if it is positive, but otherwise test again in the second period and hide a negative first outcome if the second outcome is positive. However, the DM anticipates this behaviour and based on a report of a single positive outcome she will stick with a_R, by definition her type. To convince the DM, the advisor needs to present two positive outcomes. Provided the cost is sufficiently low, the advisor will therefore continue testing after a positive outcome and induce the DM to switch if

and only if the second outcome is also positive. It is possible that the cost is so high such that the advisor is deterred from testing at all (see Figure 2.10). A single positive outcome is not enough to convince the DM and two tests are more than he is willing to pay for.

Theorem 1 [Hidden Testing vs Observable Testing]

*A **progressive** DM is strictly better off under hidden testing than observable testing if and only if* $\overline{c}^R_{01}\left(\frac{y_A}{x_A}\right) < c \leq \overline{c}^R_{02}\left(\frac{y_A}{x_A}\right)$ *(implies advisor is more reluctant than the DM and the cost is low). In this case, the advisor is also strictly better off.*

*A **moderate leaning towards conservative** DM is strictly better off under hidden testing than observable testing if and only if* $c < min\left\{\overline{c}^-_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^R_{02}\left(\frac{y_A}{x_A}\right)\right\}$ *(implies that advisor is more enthusiastic than the DM and the cost is low). In this case, the advisor is strictly worse off,*

Otherwise, the DM is never strictly better off.

There are two situations in which a DM is better off under hidden testing. First, consider the situation in which the DM is progressive and the advisor is more reluctant than the DM such that he would switch if and only if two outcomes are positive. Under observable testing, the advisor never tests and the DM never switches, if the cost is sufficiently low. By contrast, under hidden testing, the advisor tests in the first period and again in the second period if the first outcome is positive. If he discovers two positive outcomes then he reveals these to the DM, otherwise he reports a single negative outcome. The DM switches if and only if the advisor reveals two positive outcomes. If she receives a report of a single negative outcome then it could mean that the advisor has hidden a positive outcome, in which case she would prefer to switch. However, it could also mean that the advisor discovered a negative outcome on the first test, in which case she would prefer to remain. Overall, she sticks with a_R because the risk that the first test was negative is too high. Thus, the DM is better off under hidden testing, and the same applies to the advisor, because the advisor has the opportunity to gather enough evidence until he can make up his mind about whether switching is optimal. In case he is not convinced, he can prevent the DM from switching by selectively revealing negative evidence. Thus he has complete control over the DM's actions. By contrast, under observable testing, he has no control over the DM's actions and hence has no incentive to even start testing.

Furthermore, consider the situation in which the DM is moderate leaning towards conservative and

the advisor is more enthusiastic than the DM such that he would switch based on one positive and one negative outcome. Provided the cost is sufficiently low, the DM is better off under hidden testing. Under observable testing, the advisor will test in the first period but never in the second and the DM will switch if and only if the test outcome is positive. Under hidden testing, the advisor will test in the first period and again in the second period following a positive outcome and will disclose all outcomes. The DM will switch if and only if there are two positive outcomes. Although the DM is willing to switch based on one positive outcome if this is the only outcome, the DM is better off under hidden testing because she observes a second test outcome and can make her decision dependent on whether the second outcome validates the first. Under hidden testing, the advisor will test such as to have the possibility to produce two positive test outcomes because otherwise he cannot convince the DM to switch. The DM is skeptical when the advisor only reveals a single positive outcome. The DM knows that this may be the outcome of a second test after a negative first test, in which case she would like to choose a_R. It may also be the outcome of the only test, in which case she would like to choose a_S. Given that she is moderate leaning towards conservative, her preferences are such that she will remain based on the report of a single positive outcome. To overcome the DM's skepticism that he is hiding outcomes the advisor aims to find a second positive outcome. The advisor is worse off because he can convince the DM less often and invests more in testing.

The reason why the DM is better off under hidden testing in the two situations characterised in Theorem 1 are fundamentally different. To highlight this contrast, I suppose that testing is observable and compare how the DM would have to commit to future behaviour in order to achieve the payoff she receives under hidden testing. In addition, I compare whether or not she can achieve the same payoff as under hidden testing by delegating the decision making to the advisor. If the DM delegates, the advisor behaves as if he was a single player (see Section 2.3).

Proposition 1 [Commitment and Delegation]

Suppose testing is observable.

Suppose the DM is ***progressive*** *and* $\overline{c}^R_{01}\left(\frac{y_A}{x_A}\right) < c \leq \overline{c}^R_{02}\left(\frac{y_A}{x_A}\right)$ *(implies that the advisor is more reluctant than the DM and the cost is low). Then the DM achieves the same payoff as under hidden testing if i) she commits to* a_R *based on one positive and one negative outcome (i.e. commits to behave as if she was moderate) or ii) if she delegates the decision to the advisor.*

Suppose the DM is ***moderate leaning towards conservative*** *and* $c < min\left\{\overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)\right\}$ *(implies that the advisor is more enthusiastic than the DM and the cost is low). Then the DM achieves the same payoff as under hidden testing if she commits to* a_R *based on one positive outcome in a single test (i.e. commits to behave as if she was conservative), but if she delegates the decision to the advisor then she achieves a strictly lower payoff than she would achieve under hidden testing.*

If the DM had the power to commit to how she will act on the evidence revealed under observable testing then she would be at least as well off as under hidden testing. With commitment power she can provide the advisor with incentives to conduct more tests because she can influence which test outcomes are pivotal to her decision to switch.

First, consider the situation in which the DM is progressive and the advisor is more reluctant than the DM such that he would prefer a_S if and only if two outcomes are positive. The problem is that the advisor would choose a_R after one positive outcome, but the DM would choose a_S. The advisor cannot influence the DM's decision by conducting a second test. This problem does not arise if the DM were moderate. Hence, by committing to act as if she were moderate the DM can provide the advisor with incentives to test.

Second, consider the situation in which the DM is moderate leaning towards conservative and the advisor is more enthusiastic than the DM such that he would choose a_S based on one positive and one negative outcome. The problem is that the advisor stops testing after the first period if he has discovered a positive outcome because this is sufficient to convince the DM to switch. However, the DM would value to learn whether there is stronger evidence in favour of the hypothesis. If the DM behaves as if she were conservative and switches if and only if there are two positive outcomes, then the advisor would continue testing after one positive outcome.

In the first situation the DM adjusts her behaviour to be more in line with what the advisor would like her to do, whereas in the second situation the DM adjusts her behaviour to be less in line with what the advisor would do. This difference between the two situations is also reflected in the DM's benefit from delegation. In the first situation, the progressive DM benefits from delegation and achieves the same outcome as under hidden testing. Hidden testing has given the advisor additional freedom which means he has control over the decision making. The DM benefits when the advisor has such control, because only if the advisor has control over how the information is used will he be willing to test at all.

By contrast, in the second situation, the moderate leaning towards conservative DM is strictly worse off when she delegates to the advisor. The advisor tests in order to convince the DM to switch. If he had control over the decision making he would require no test at all or fewer positive outcomes in order to be convinced. Hidden testing allows him more freedom, but this only causes the DM to be more skeptical towards his report. Since it is more difficult to convince the DM, the advisor is more motivated to test.

So far I have focused on situations in which the DM is better off under hidden testing than under observable testing. It is worth noting that it is also possible that both the DM and the advisor are worse off under hidden testing.

Proposition 2 [Hidden Testing Pareto Dominated]

*Suppose the DM is **moderate leaning towards conservative.** Then both the DM and the advisor are strictly worse off under hidden testing than observable testing if and only if* $\overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right) < c < \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$ *(implies that the advisor is more enthusiastic than the DM and the cost is intermediate).*

Suppose the DM is moderate leaning towards conservative and the advisor is so enthusiastic that he would switch based on one positive and one negative outcome. If costs fall into some intermediate range, the advisor does not test at all under hidden testing, whereas he would have tested in the first period if testing were observable. On the one hand, the cost is so low that the DM is skeptical when she observes one positive test outcome in a single test because the advisor may have conducted two tests and not disclosed what he considers unfavourable evidence. On the other hand, the cost is so high that the advisor does not find it worthwhile to follow a testing procedure such that he can discover two positive outcomes to convince the DM. Thus, he does not test at all. Consequently, the DM is worse off because she never learns valuable evidence about the hypothesis and the advisor is worse off because he can never convince the DM to switch.

2.6 Choosing an Advisor

The extent to which the advisor's preferences differ from the DM's preferences influences the advisor's optimal testing procedure. Are there situations in which the DM prefers a biased to an unbiased advisor? When would she prefer an advisor who is more reluctant towards switching than herself, and when would

she prefer an advisor who is more enthusiastic about switching than herself? I first consider this question when testing is observable and then when testing is hidden.

When testing is observable the DM faces the following trade-off: if an unbiased advisor is motivated to test then he will follow the same testing procedure the DM would have followed. However, testing is costly and therefore the unbiased advisor may not conduct all the tests the DM would like him to conduct. A biased advisor may not follow the optimal testing procedure, but he may be willing to test when the unbiased advisor is not willing to test.

Proposition 3 [Biased vs Unbiased Advisor - Observable Testing]

Suppose testing is observable and that the loss x from choosing a_R if the hypothesis is true is the same for the advisor as for the DM, i.e. $x_A = x_{DM}$.

Then the DM is weakly better off the lower y_A if

- *the DM is progressive, or*
- *the DM is moderate and $c > \widehat{c}$, where $\widehat{c} = \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$, or*
- *the DM is conservative.*

If the DM is moderate and $c < \widehat{c}$, then her preferences over y_A are single peaked and she may be strictly better off with a more reluctant advisor ($y_A > y_{DM}$).

Consider a situation in which the DM would choose a_R if the advisor were to refrain from testing, but if the advisor were to test then some test outcome would lead the DM to choose a_S. If testing was costless, an unbiased advisor would always test in such a situation because the test would improve the DM's decision, but it may be that the cost is too high for him to find testing worthwhile. Then a more reluctant advisor would also not test, but a more enthusiastic advisor may still be willing to test, because a more enthusiastic advisor reduces his loss by more than an unbiased advisor if he finds evidence in favour of the hypothesis and induces the DM to choose a_S.[6]

Now consider a situation in which the DM would choose a_S if the advisor were to refrain from testing, but if the advisor were to test then some test outcome would lead the DM to choose a_R. Such a situation

[6]This is because the more enthusiastic advisor incurs a lower loss y_A than the unbiased advisor if the DM chooses a_S and the hypothesis is false. Therefore, the more enthusiastic advisor finds it relatively more important to avoid loss x_A incurred if the DM chooses a_R and the hypothesis is true.

can only arise if the DM is moderate and the advisor's preferences and the cost are such that the advisor tests in the first period but not in the second period. If the first outcome is positive then the DM would choose a_S without a further test, but if a second test was negative she would revert back to a_R. A more reluctant advisor may still be willing to conduct an additional test following a positive outcome. This is because a more reluctant advisor incurs a higher loss y_A than an unbiased advisor if the DM chooses a_S and the hypothesis is false.

Proposition 4 [Biased vs Unbiased Advisor - Hidden Testing]

Suppose testing is hidden and that the loss x from choosing a_R if the hypothesis is true is the same for the advisor as for the DM, i.e. $x_A = x_{DM}$.

In contrast to the case of observable testing:

- *if $c > \widehat{c}$ and the DM is moderate (either leaning towards progressive or conservative) then the DM may be strictly worse off as y_A is lower, and*
- *if $c \leq \widehat{c}$ and the DM is moderate leaning towards conservative then the DM's preferences over y_A are not single-peaked.*

Otherwise, the statements of Proposition 3 also apply under hidden testing.

When testing is hidden then an unbiased advisor would disclose all outcomes, but a biased advisor may hide some outcomes, which is not desirable for the DM. However, a biased advisor may also follow a different testing procedure under hidden testing than under observable testing, as Section 2.5 has shown. Therefore, a biased advisor is not necessarily less attractive for the DM when testing is hidden rather than observable.

Suppose the DM is moderate and the cost is so high that under observable testing a more enthusiastic advisor would have been weakly preferred. Under hidden testing the more enthusiastic advisor is so motivated to convince the DM to switch that he is tempted to hide a negative outcome on the first test if the second test is positive. This is not desirable for the DM, who would choose a_R after a single negative outcome. If the DM is leaning towards progressive, she may end up choosing a_S based on weaker evidence. If the DM is leaning towards conservative, the advisor may not test at all because he cannot convince the DM to choose a_S with the report of a single positive outcome and finds it too costly

to produce two positive outcomes.

Now suppose the DM is moderate leaning towards conservative and the the cost is so low that under observable testing the DM would never have been strictly better off with a more enthusiastic advisor. Under hidden testing, it is possible that a more enthusiastic advisor tests more often than an unbiased advisor. The unbiased advisor may only test in the first period and never in the second. By contrast, a more enthusiastic advisor is motivated to test again following a positive outcome. This is because the DM is skeptical that he may be hiding a negative outcome if he reports one positive outcome in a single test. Unlike an unbiased advisor, the more enthusiastic advisor has to overcome this skepticism. Therefore, either a more enthusiastic or a more reluctant advisor would ensure that the DM will learn whether or not two outcomes were positive.

2.7 Related Literature

There is a large literature analysing a decision maker (DM) receiving advice from an imperfectly informed advisor. A part of this literature considers verifiable disclosure when the advisor's information is exogenously given, e.g. Milgrom [1981], Fishman and Hagerty [1990]. The work most closely related to this paper considers a situation in which the advisor first acquires information and then selectively discloses what he has learned. Unlike other papers in the literature, I focus on the contrast between hidden and observable testing when the advisor's preferences over actions are state-dependent.

2.7.1 Non-sequential Testing

A part of the literature also studies how hidden testing can motivate the advisor to acquire more evidence, because the DM is skeptical that the advisor is hiding contradicting evidence.

2.7.1.1 Che and Kartik (2009)

Che and Kartik [2009]'s main focus is how information acquisition and verifiable disclosure are affected when the advisor's prior belief differs from the DM's and what this implies for the DM's choice of advisor. In addition, they contrast their results with those derived when the advisor differs in his preferences instead of his prior belief, and they also study how difference in preferences and prior beliefs interact.

In their framework, the action and state space are continuous, and the DM's loss is captured by a

quadratic function of the difference between his action and the true state. The advisor's loss is also quadratic. When the advisor has different preferences than the DM, his ideal action given the state is either higher or lower than the DM's ideal action. The extent to which his ideal action is different measures the size of his bias. This implies that even under certainty their preferences over actions differ. To gather evidence, the advisor can privately exert costly effort which increases the probability that he will receive an informative signal. The signal realisations are normally distributed with mean equal to the true state, and the advisor can at most observe one signal. Neither the effort level nor the signal realisation are observed by the DM. The advisor then chooses whether or not to disclose the signal, and the DM chooses the action.

I will focus on the situation in which both players have the same prior belief but different preferences. Che and Kartik [2009] show that if the signal realisation were observable then the DM would be indifferent between a biased and an unbiased advisor. The advisor's bias has no impact on how much effort he exerts because with or without a test a biased advisor expects the DM's action to differ from his ideal action by the size of the bias. However, the authors find that if the signal and effort choice are not observable to the DM, then the advisor has more incentives to invest the larger his bias. This is because, when a biased advisor does not disclose a signal, the DM cannot tell whether he did not observe a signal or whether he is hiding a signal outcome that would induce the DM to choose an action which is less desirable from the advisor's perspective. Therefore, whenever a biased advisor does not report a signal, the DM adjusts her action in the direction less desired by the advisor. By contrast, if an unbiased advisor does not report a signal the DM anticipates that this is because he did not observe a signal and chooses the same action she would have chosen as a default. Since a biased advisor expects to be worse off when he does not generate a signal, he has more incentives to invest than an unbiased advisor. Overall, the DM faces a trade-off when choosing how much the advisor's preferences should differ from hers: a more biased advisor exerts more effort, but hides the signal for a larger range of outcomes. The authors show that locally the DM prefers an unbiased advisor to one who is slightly biased, but it is possible that a more biased advisor is preferred to an unbiased advisor. They also find that it is never optimal for the DM to delegate the decision to a biased advisor.

My paper differs from Che and Kartik [2009] in important aspects of the modelling framework. First, I assume that the action and state space are binary and that both players agree on which choice is optimal under certainty, and second I assume that testing is sequential. This gives rise to different findings than

in Che and Kartik [2009]. First, even when testing is observable, a biased advisor may be preferred to an unbiased advisor (Proposition 3). Since the action space is binary, a test outcome can either induce the DM to stick with the action she would have chosen as a default or it can induce her to change to the alternative action. How valuable it is for an advisor to induce the DM to change to the alternative action depends on his bias. Hence, a biased advisor may benefit more than an unbiased advisor from conducting a test. Second, under hidden testing, Che and Kartik [2009] find that for some intermediate range of the size of the bias, the DM prefers a biased advisor to an unbiased advisor. I find that not only the size of the bias matters but also the direction of the bias. In my paper, for some range of preferences, the DM may benefit from an advisor who is more enthusiastic about switching than herself, but for another range of preferences, the DM may benefit from an advisor who is more reluctant to switch than herself (Proposition 4). The reason for why she can benefit from a more enthusiastic advisor is not identical to the reason for why a DM may prefer a biased advisor in Che and Kartik [2009], but it has a similar flavour. In my model, the more enthusiastic advisor is tempted to continue testing in order to discover positive evidence and hide negative evidence. Since the DM anticipates this behaviour, she suspects that the advisor is not revealing all outcomes. In order to overcome the DM's skepticism, the advisor tests more than under observable testing. In particular, he conducts a further test after a positive outcome to validate the outcome of the first test rather than after a negative outcome to contradict the finding of the first test. The reason for why a DM may benefit from a more reluctant advisor does not appear in Che and Kartik [2009]. In this case, I also find that delegating the decision making to the advisor can make the DM better off.

2.7.1.2 Henry (2009)

Henry [2009] considers the impact of hidden testing when the advisor has to commit ex-ante to how much evidence he acquires rather than testing sequentially. He finds that the DM is always better off under hidden testing. In his framework, the reason is that the advisor will test more because has more discretion of what to report, but since, in equilibrium, the DM knows how much evidence the advisor has collected the advisor will not find it optimal to hide any evidence.

In my paper, I allow the advisor to test sequentially to allow additional scope for strategic influence. The advisor can decide whether to stop or continue testing on the basis of the outcome of an earlier test. This assumption is realistic in many scientific research environments in which experiments are first

evaluated before further experiments are started. I find that a more enthusiastic advisor is motivated to present stronger evidence to overcome the DM's skepticism, thereby demonstrating that existing results derived under the assumption of simultaneous testing are robust to allowing for sequential testing.

Importantly, in contrast to Henry [2009], I find an additional motivation for why an advisor may collect more evidence under hidden testing. If the advisor is more reluctant to switch from the default action than the DM and testing is observable, then he may not have any incentive to acquire evidence, because he may risk that the evidence is too weak for the advisor himself to be convinced, but strong enough for the DM to be convinced to switch. By contrast, under hidden testing, the advisor can first privately explore whether or not the evidence is sufficient to convince himself that switching is optimal without fearing any adverse consequences for the DM's choice, because he then still has the option to withhold this evidence and the DM will stick with his default choice.

2.7.2 Sequential Testing

In contrast to existing work on sequential testing, the focus of my paper is on how the DM is affected if testing is hidden rather than observable.

2.7.2.1 Felgenhauer and Schulte (2014)

Felgenhauer and Schulte [2014] study a persuasion game between an advisor and a DM and focus on how persuasion depends on the advisor's cost of testing, when the advisor can privately acquire as many signals as he likes before sending a report to the DM.

They assume that the DM can choose one of two actions and her loss from choosing an action depends on the state. By contrast, the advisor incurs a state-independent benefit whenever the DM changes her action from the default. The advisor can sequentially acquire an arbitrary number of private binary signals, each at a fixed cost and generated using the same technology. Then he can verifiably disclose a selection of signals to the DM. The DM can neither observe how often the advisor has tested nor can she make any inference about the number of tests if the advisor reports no outcomes, i.e. the timing of the report conveys no information.

Felgenhauer and Schulte [2014] focus on equilibria in which the DM's strategy is to stick with the default action unless the advisor reveals more than a certain number of positive test outcomes. In such an equilibrium, the advisor will reveal all positive outcomes and hide all negative outcomes. Once the

advisor has found enough positive outcomes to convince the DM then he will reveal these outcomes. However, the advisor may stop even if he has insufficient evidence to convince the DM and report no findings to the DM. This happens if the evidence required up to this point leads the advisor to become too pessimistic regarding the chances of persuading the DM. There exists a range of equilibria with different threshold levels of evidence for the DM. Felgenhauer and Schulte [2014] find that the smallest and largest thresholds increase as the advisor's cost of testing decreases. In an extension, they consider a situation in which only a limited number of signals can be conveyed by the advisor. If the cost is sufficiently low, then no testing arises in equilibrium.

An important difference in my framework is that I allow the advisor's preferences to depend on the state. In my model, the special case in which $\frac{y_A}{x_A} = 0$ corresponds to an advisor who would want the DM to choose a_S whatever the state because he never suffers a loss from choosing a_S if the hypothesis is false. Assuming $\frac{y_A}{x_A} = 0$ and hidden testing, then my findings are in line with those of Felgenhauer and Schulte [2014] (Lemma 3). In particular, if the DM is moderate leaning towards conservative then she requires two outcomes to switch if the cost is low and one outcome if the cost is higher (and no equilibrium with testing exists if the cost is very high). Because I analyse a situation in which the advisor can draw at most two signals, the DM does not face an inference problem about which test outcomes the advisor has observed whenever he reveals two outcomes. While I am imposing restrictions on the number of tests, I allow for a wider range of advisor's preferences. While Felgenhauer and Schulte [2014] find that the cost of testing influences how credible a given report is, I show that not only the cost but also the divergence in preferences matters. A crucial difference between my paper and theirs is that I show that, from the DM's point of view, hidden testing may be preferable to observable testing.

2.7.2.2 Gerardi and Yariv (2008)

Gerardi and Yariv [2008] study how a DM can optimally elicit advice from a committee of advisors when she can choose how much each advisor's preferences diverge from her own.

The DM has to take a binary decision and favours a different action depending on which of two states of the world is realised. The state of the world is unknown and she can consult an arbitrary number of advisors for advice. Advisors can invest to become informed. The more an advisor invests at a private cost, the more accurate is the binary signal he receives.[7] While an advisor's effort level is not observed

[7]The accuracy of the signal is independent of the state.

by the DM, the signal realisation is public knowledge. The DM has no commitment power and hence any proposed mechanism needs to be incentive-compatible for the DM to follow, given the advisors' strategies and the information available to the DM at the time she takes a decision.

Gerardi and Yariv [2008] show that if the DM consults an advisor at all, then she follows a sequential mechanism in which she consults two advisors: The DM first consults one advisor. If his signal is negative, i.e. if the signal realisation confirms her prior belief, then she will stick with her default choice. If the signal is positive, then she will consult the other advisor and change her action if and only if his signal realisation is positive. The DM chooses the two advisors such that their preferences are extreme and opposed to her own preferences, i.e. each advisor would choose the alternative action irrespective of the evidence gathered.

Gerardi and Yariv [2008] find that the optimal advisor is one who is extremely enthusiastic about choosing the alternative action. I also find that the DM may prefer a more enthusiastic advisor to an unbiased advisor. The underlying reasoning is the same in both papers: a more biased advisor may gain more by inducing the DM to switch from her default choice and hence he is more motivated to test. However, I also find that the DM may sometimes prefer a more reluctant advisor to an unbiased advisor (Proposition 3). Suppose the unbiased advisor stops testing after the first period and the (moderate) DM switches if the outcome is positive. Then a more reluctant advisor may test again following a positive outcome. This is because the DM will choose a_S if she does not observe a further test and therefore the preferred advisor is one who is opposed to this choice.

2.8 Conclusion

This paper studies a DM who has to decide whether or not to change her action based on the evidence gathered by an advisor. The advisor agrees with the DM that she should change her action if and only if a certain hypothesis is true, yet the advisor may incur a relatively larger or a relatively smaller loss than the DM if the hypothesis is falsely accepted. Over the course of two periods, the advisor can sequentially acquire informative signals, each at a fixed cost, using an exogenously given testing technology. I contrast the situation in which testing is publicly observed with the situation in which testing is privately observed and the advisor can selectively disclose outcomes.

I find that there are two kinds of situations in which the DM is strictly better off if testing is hidden

rather than observable. The first situation arises when the advisor is more reluctant to change action than the DM because he would incur a relatively higher loss from falsely accepting than the DM. Under observable testing, the advisor will not test because the DM will then stick with the default action for certain, whereas from the advisor's point of view testing could seduce the DM to switch prematurely. Under hidden testing the advisor has more influence over the DM's decision because he can selectively hide outcomes. Due to this additional scope for influence, the advisor has incentives to acquire evidence at all.

The second situation arises when the advisor is more enthusiastic about changing action than the DM because he would incur a relatively higher loss from falsely rejecting than the DM. The advisor will stop testing once he has convinced the DM to switch action. Under hidden testing the DM requires stronger evidence to switch action than under observable testing. This is because under hidden testing the DM is skeptical that the advisor is hiding contradicting evidence and hence wants to see that additional tests confirm earlier findings.

These insights have implications for designing the appropriate research environment. In many circumstances, researchers have the freedom to decide when they stop accumulating evidence. Often the public or interested party who base their decisions on the researcher's findings have different preferences than the researcher himself. For example, the researcher may hesitate to publish findings which are scientifically true because he dreads a large loss in reputation if later on his findings are refuted. The public on the other hand may be eager to adapt in case some novel hypothesis turns out to be true, e.g. they may want to change their behaviour if there is any indication that their current conduct could be harmful. To reduce the researcher's scope for manipulation, institutions may be able to monitor the researcher more closely. However, my findings show that the optimal policy may be to give the researcher even more scope for manipulation and allow him to conduct his research in private and selectively disclose findings. This is particularly valuable if the cost of acquiring evidence is low. In some cases, a researcher is then forced to produce stronger evidence if he cannot credibly deny having withheld contradicting evidence. In other cases, a researcher will have an incentive to investigating a certain hypothesis only if he can have some control over confirming his findings with additional tests before releasing his results.

In addition, this paper shows that under observable testing a DM should choose an advisor who opposes his default choice if the cost of testing is high. However, if the DM values having additional confirmation of early positive findings and if the cost is low then she might be better off choosing an

advisor who is more reluctant to deviate from the default choice than the DM herself. Under hidden testing, not only a more reluctant but also a more enthusiastic advisor may be willing to check whether earlier findings are robust to additional tests when an unbiased advisor would not conduct such robustness checks.

In future work, it would be valuable to analyse how the difference between observable and hidden testing affects payoffs if the DM could consult several advisors.

2.9 Appendix

Definition A1:

Denote by $\widetilde{h}_2$ the unordered history of test outcomes. Any history of test outcomes h_2 maps into an unordered history $\widetilde{h}_2$. The unordered history only contains information about the number of positive and the number of negative signals, but not their order, i.e. $\widetilde{h}_2 \in \{(++),(+-),(--),(+),(-),(\emptyset)\}$. For example, $h_2 = (\emptyset+)$ maps into $\widetilde{h}_2 = (+)$.

2.9.1 Lemma 1

For a single player, the number of positive and the number of negative test outcomes is the only information relevant for his decision. I will characterise his decision rule by the set Z of unordered histories of test outcomes (Def A1) such that the player chooses a_S if and only if $\widetilde{h}_2 \in Z$. I will divide the range of preferences $\frac{y}{x}$ into sections such that the player's decision rule Z is constant within each section and different across sections. Then I will solve for the optimal testing strategy, denoted by T, using backward induction. Throughout, I will use $Pr\left(\omega^i|x\right)$ to denote the posterior belief that $\omega = \omega^i$ conditional on $\widetilde{h}_2 = (x)$.

2.9.1.1 $Z = \{(++),(+-),(--),(+),(-),(\emptyset)\}$.

This is the decision rule if and only if

$$\frac{y}{x} \leq \frac{q(1-p_1)^2}{(1-q)p_2} = \frac{Pr\left(\omega^T|--\right)}{Pr\left(\omega^F|--\right)}.$$

Since the advisor will choose a_S for any $\widetilde{h}_2$, he will never test:

$$T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right).$$

2.9.1.2 $Z = \{(++),(+-),(+),(-),(\emptyset)\}$.

This is the decision rule if and only if

$$\frac{q(1-p_1)^2}{(1-q)p_2} \leq \frac{y}{x} < \frac{q(1-p_1)}{(1-q)p_2} = \frac{Pr\left(\omega^T|-\right)}{Pr\left(\omega^F|-\right)}.$$

Second Period:

$t_2^+ = N$: Given $h_1 = (+)$, he would choose a_S regardless of the outcome of the second test.

$t_2^- = Y$ if and only if $c < \underline{c}^{-}\left(\frac{y}{x}\right)$: Given $h_1 = (-)$, he would choose a_S without a further test, but if and only if the second test was negative then he would choose a_R. Hence, $t_2^- = Y$ if and only if

$$\begin{aligned} & Pr\left(\omega^T|-\right)(1-p_1) & & \\ +\left(1-Pr\left(\omega^T|-\right)\right)\frac{y}{x}(1-p_2)+c & < & \left(1-Pr\left(\omega^T|-\right)\right)\frac{y}{x} & \\ & \Leftrightarrow & & \\ c & < & \frac{p_2(1-q)\frac{y}{x}-(1-p_1)^2 q}{p_2(1-q)+(1-p_1)q} \equiv \underline{c}^{-}\left(\frac{y}{x}\right). & (2.1) \end{aligned}$$

$t_2^{\emptyset} = N$: Given $h_1 = (\emptyset)$, he would choose a_S regardless of the outcome of the test.

First Period:

If $c > \underline{c}^{-}\left(\frac{y}{x}\right)$ then he would neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right).$$

If $c < \underline{c}^{-}\left(\frac{y}{x}\right)$, he would choose a_S without a further test, but if and only if both tests were negative then

he would choose a_R. Hence $t_1 = Y$ if and only if

$$\begin{aligned} & Pr\left(\omega^T\right)(1-p_1)^2 \\ & +\left(1-Pr\left(\omega^T\right)\right)\frac{y}{x}(1-p_2) \\ & +c\left(1+p_2(1-q)+(1-p_1)q\right) \quad < \quad \left(1-Pr\left(\omega^T\right)\right)\frac{y}{x} \\ & \Leftrightarrow \\ & c \quad < \quad \frac{p_2(1-q)\frac{y}{x}-(1-p_1)^2 q}{1+p_2(1-q)+(1-p_1)q} \equiv \overline{c}_{02}^{S}\left(\frac{y}{x}\right). \end{aligned} \tag{2.2}$$

Clearly, $\overline{c}_{02}^{S}\left(\frac{y}{x}\right) < \underline{c}^{-}\left(\frac{y}{x}\right)$.

If $\overline{c}_{02}^{S}\left(\frac{y}{x}\right) < c < \underline{c}^{-}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{02}^{S}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = N\right).$$

2.9.1.3 $Z = \{(++),(+-),(+),(\emptyset)\}$.

This is the decision rule if and only if

$$\frac{q(1-p_1)}{(1-q)p_2} \leq \frac{y}{x} < min\left\{\frac{q}{1-q}, \frac{q(p_1+p_1(1-p_1))}{(1-q)(1-p_2+p_2(1-p_2))}\right\}$$

where

$$\frac{q(p_1+p_1(1-p_1))}{(1-q)(1-p_2+p_2(1-p_2))} = \frac{Pr\left(\omega^T|+-\right)}{Pr\left(\omega^F|+-\right)}.$$

Second Period:

$t_2^{+} = N$: Given $h_1 = (+)$, he would choose a_S regardless of the outcome of the second test.

$t_2^{-} = Y$ if and only if $c < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$: Given $h_1 = (-)$, he would choose a_R without a further test, but if and

only if the second test was positive then he would choose a_S. Hence, $t_2^- = Y$ if and only if

$$\begin{aligned} & Pr\left(\omega^T|-\right)(1-p_1) & & \\ +\left(1-Pr\left(\omega^T|-\right)\right)\frac{y}{x}(1-p_2)+c & < & Pr\left(\omega^T|-\right) & \\ & \Leftrightarrow & & \\ c & < & \frac{p_1(1-p_1)q-p_2(1-p_2)(1-q)\frac{y}{x}}{p_2(1-q)+(1-p_1)q} \equiv \overline{c}_{12}^{-}\left(\frac{y}{x}\right). & (2.3) \end{aligned}$$

$t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$: Given $h_1 = (\emptyset)$, he would choose a_S without a further test, but if and only if the test was negative then he would choose a_R. Hence $t_2^{\emptyset} = Y$ if and only if

$$\begin{aligned} & Pr\left(\omega^T\right)(1-p_1) & & \\ +\left(1-Pr\left(\omega^T\right)\right)\frac{y}{x}(1-p_2)+c & < & \left(1-Pr\left(\omega^T\right)\right)\frac{y}{x} & \\ & \Leftrightarrow & & \\ c & < & (1-q)\frac{y}{x}p_2 & \\ & & -q(1-p_1) \equiv \overline{c}_{01}^{S}\left(\frac{y}{x}\right). & (2.4) \end{aligned}$$

First period:

If $c > max\left\{\overline{c}_{01}^{S}\left(\frac{y}{x}\right), \overline{c}_{12}^{-}\left(\frac{y}{x}\right)\right\}$ then he would neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then $t_1 = Y$: He tests because he would also start testing in the second period.

If $\overline{c}_{12}^{-}\left(\frac{y}{x}\right) < c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = Y\right).$$

If $c < \overline{c}_{12}^{-}\left(\frac{y}{x}\right) < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^+ = N, t_2^- = Y, t_2^{\emptyset} = Y\right).$$

If $\overline{c}_{01}^{S}\left(\frac{y}{x}\right) < c < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$ then he would choose a_S without a further tests, but if he tests in the first period and then again in the second period following a negative outcome, he would choose a_R if and only if two outcomes are negative. Hence $t_1 = Y$ if and only if $c < \overline{c}_{02}^{S}\left(\frac{y}{x}\right)$ (see equation (2.2)). If $\overline{c}_{01}^{S}\left(\frac{y}{x}\right) < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$ then $\overline{c}_{01}^{S}\left(\frac{y}{x}\right) < \overline{c}_{02}^{S}\left(\frac{y}{x}\right) < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$.

If $\overline{c}_{01}^{S}\left(\frac{y}{x}\right) < c < \overline{c}_{02}^{S}\left(\frac{y}{x}\right) < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = N\right).$$

If $\overline{c}_{01}^{S}\left(\frac{y}{x}\right) < \overline{c}_{02}^{S}\left(\frac{y}{x}\right) < c < \overline{c}_{12}^{-}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = N\right).$$

2.9.1.4 $Z = \{(++),(+)\}$.

This is the decision rule if and only if

$$max\left\{\frac{q}{1-q}, \frac{q\left(p_1 + p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(1-p_2+p_2\left(1-p_2\right)\right)}\right\} \leq \frac{y}{x} < \frac{qp_1}{\left(1-q\right)\left(1-p_2\right)}$$

where

$$\frac{qp_1}{\left(1-q\right)\left(1-p_2\right)} = \frac{Pr\left(\omega^T|+\right)}{Pr\left(\omega^F|+\right)}.$$

Second Period:

$t_2^{+} = Y$ if and only if $c < \overline{c}_{12}^{+}\left(\frac{y}{x}\right)$: Given $h_1 = (+)$, he would choose a_S without a further test, but if and only if the second test was negative then he would choose a_R. Hence, $t_2^{+} = Y$ if and only if

$$\begin{aligned} & Pr\left(\omega^T|+\right)\left(1-p_1\right) & & \\ +\left(1-Pr\left(\omega^T|+\right)\right)\frac{y}{x}\left(1-p_2\right)+c & < & \left(1-Pr\left(\omega^T|+\right)\right)\frac{y}{x} \\ & \Leftrightarrow & \\ c & < & \frac{p_2\left(1-p_2\right)\left(1-q\right)\frac{y}{x} - p_1\left(1-p_1\right)q}{\left(1-p_2\right)\left(1-q\right)+p_1 q} \equiv \overline{c}_{12}^{+}\left(\frac{y}{x}\right). \end{aligned} \quad (2.5)$$

$t_2^{-} = N$: Given $h_1 = (-)$, he would choose a_R regardless of the outcome of the second test.

$t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$: Given $h_1 = (\emptyset)$, he would choose a_R without a further test, but if and

only if the test was positive then he would choose a_S. Hence $t_2^{\emptyset} = Y$ if and only if

$$\begin{aligned} Pr\left(\omega^T\right)(1-p_1) & \\ +\left(1-Pr\left(\omega^T\right)\right)\frac{y}{x}(1-p_2)+c &< Pr\left(\omega^T\right) \\ \Leftrightarrow & \\ c &< qp_1-(1-q)(1-p_2)\frac{y}{x} \equiv \overline{c}_{01}^R\left(\frac{y}{x}\right). \end{aligned} \tag{2.6}$$

First period:

If $c > max\left\{\overline{c}_{01}^R\left(\frac{y}{x}\right), \overline{c}_{12}^+\left(\frac{y}{x}\right)\right\}$ then he would neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^R\left(\frac{y}{x}\right)$ then $t_1 = Y$: He will test because he would also start testing in the second period.

If $\overline{c}_{12}^+\left(\frac{y}{x}\right) < c < \overline{c}_{01}^R\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = Y\right).$$

If $c < \overline{c}_{12}^+\left(\frac{y}{x}\right) < \overline{c}_{01}^R\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^+ = Y, t_2^- = N, t_2^{\emptyset} = Y\right).$$

If $\overline{c}_{01}^R\left(\frac{y}{x}\right) < c < \overline{c}_{12}^+\left(\frac{y}{x}\right)$ then the DM would choose a_R without a further test, but if the advisor tests in the first period and then again in the second period following a positive outcome, she would choose a_S if both tests were positive. Hence $t_1 = Y$ if and only if

$$\begin{aligned} Pr\left(\omega^T\right)(1-p_1) & \\ +\left(1-Pr\left(\omega^T\right)\right)\frac{y}{x}(1-p_2)^2 & \\ +c\left(1+(1-p_2)(1-q)+p_1q\right) &< Pr\left(\omega^T\right) \\ \Leftrightarrow & \\ c &< \frac{p_1q-(1-p_2)^2(1-q)\frac{y}{x}}{1+(1-p_2)(1-q)+p_1q} \equiv \overline{c}_{02}^R\left(\frac{y}{x}\right). \end{aligned} \tag{2.7}$$

If $\overline{c}_{01}^{R}\left(\frac{y}{x}\right) < \overline{c}_{12}^{+}\left(\frac{y}{x}\right)$ then $\overline{c}_{01}^{R}\left(\frac{y}{x}\right) < \overline{c}_{02}^{R}\left(\frac{y}{x}\right) < \overline{c}_{12}^{+}\left(\frac{y}{x}\right)$.

If $\overline{c}_{01}^{R}\left(\frac{y}{x}\right) < c < \overline{c}_{02}^{R}\left(\frac{y}{x}\right) < \widetilde{c}^{+}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $\overline{c}_{01}^{R}\left(\frac{y}{x}\right) < \overline{c}_{02}^{R}\left(\frac{y}{x}\right) < c < \widetilde{c}^{+}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = N, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

2.9.1.5 $Z = \{(++)\}$.

This is the decision rule if and only if

$$\frac{qp_1}{(1-q)(1-p_2)} \leq \frac{y}{x} < \frac{qp_1}{(1-q)(1-p_2)^2} = \frac{Pr\left(\omega^T|++\right)}{Pr\left(\omega^F|++\right)}.$$

Second Period:

$t_2^{+} = Y$ if and only if $c < \widetilde{c}^{+}\left(\frac{y}{x}\right)$: Given $h_1 = (+)$, he would choose a_R without a further test, but if and only if the second test was positive then he would choose a_S. Hence, $t_2^{+} = Y$ if and only if

$$\begin{aligned} Pr\left(\omega^T|+\right)(1-p_1) & \\ +\left(1-Pr\left(\omega^T|+\right)\right)\frac{y}{x}(1-p_2)+c \quad &< \quad Pr\left(\omega^T|+\right) \\ &\Leftrightarrow \\ c \quad &< \quad \frac{p_1 q-(1-p_2)^2(1-q)\frac{y}{x}}{(1-p_2)(1-q)+p_1 q} \equiv \widetilde{c}^{+}\left(\frac{y}{x}\right). \end{aligned} \tag{2.8}$$

$t_2^{-} = N$: Given $h_1 = (-)$, he would choose a_R regardless of the outcome of the second test.

$t_2^{\emptyset} = N$: Given $h_1 = (\emptyset)$ he would choose a_R regardless of the outcome of the test.

First Period:

If $c > \widetilde{c}^{+}\left(\frac{y}{x}\right)$ he would neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $c < \widetilde{c}^{+}\left(\frac{y_A}{x_A}\right)$, he would choose a_R without a further test, but if and only if both tests were positive then he would choose a_S. Hence $t_1 = Y$ if and only if $c < \overline{c}^{R}_{02}\left(\frac{y}{x}\right)$ (see equation (2.7)).

If $c < \overline{c}^{R}_{02}\left(\frac{y}{x}\right) < \widetilde{c}^{+}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $\overline{c}^{R}_{02}\left(\frac{y}{x}\right) < c < \widetilde{c}^{+}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = N, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

2.9.1.6 Z is empty.

This is the decision rule if and only if

$$\frac{y}{x} \geq \frac{qp_1}{(1-q)(1-p_2)^2} = \frac{Pr\left(\omega^T | ++\right)}{Pr\left(\omega^F | ++\right)}$$

Since he will choose a_R for any $\widetilde{h}_2$, he will never test:

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

2.9.1.7 $Z = \{(++),(+),(+-)\}$

This is the decision rule if and only if $p_1 > p_2$ and

$$\frac{q}{1-q} \leq \frac{y}{x} < \frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))} = \frac{Pr\left(\omega^T | +-\right)}{Pr\left(\omega^F | +-\right)}.$$

Second Period:

$t_2^{+} = N$: Given $h_1 = (+)$, he would choose a_S regardless of the outcome of the second test.

$t_2^{-} = Y$ if and only if $c < \overline{c}^{-}_{12}\left(\frac{y}{x}\right)$: Given $h_1 = (-)$, he would choose a_R without a further test, but if and only if the second test was positive then he would choose a_S. Hence, $t_2^{-} = Y$ if and only if $c < \overline{c}^{-}_{12}\left(\frac{y}{x}\right)$ (see equation (2.3)).

$t_2^{\emptyset} = Y$ if and only if $c < \overline{c}^{R}_{01}\left(\frac{y}{x}\right)$: Given $h_1 = (\emptyset)$, he would choose a_R without a further test, but if and only if the test was positive then he would choose a_S. Hence $t_2^{\emptyset} = Y$ if and only if $c < \overline{c}^{R}_{01}\left(\frac{y}{x}\right)$ (see equation

(2.6)).

We have $\overline{c}_{12}^{-}\left(\frac{y}{x}\right) < \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$ since $Pr\left(\omega^T\right) > Pr\left(\omega^T|-\right)$.

First period:

If $c > \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$ then he would neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$ then $t_1 = Y$: He would test because he would also start testing in the second period.

If $\overline{c}_{12}^{-}\left(\frac{y}{x}\right) < c < \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

If $c < \overline{c}_{12}^{-}\left(\frac{y}{x}\right) < \overline{c}_{01}^{R}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = Y\right).$$

2.9.1.8 $Z = \{(++),(+),(\emptyset)\}$

This is the decision rule if and only if $p_2 > p_1$ and

$$\frac{q\left(p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(p_2\left(1-p_2\right)\right)} \leq \frac{y}{x} < \frac{q}{1-q}.$$

Second Period:

$t_2^{+} = Y$ if and only if $c < \overline{c}_{12}^{+}\left(\frac{y}{x}\right)$: Given $h_1 = (+)$, he would choose a_S without a further test, but if and only if the second test was negative then he would choose a_R. Hence, $t_2^{+} = Y$ if and only if $c < \overline{c}_{12}^{+}\left(\frac{y}{x}\right)$ (see equation (2.5)).

$t_2^{-} = N$: Given $h_1 = (-)$, he would choose a_R regardless of the outcome of the second test.

$t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$: Given $h_1 = (\emptyset)$, he would choose a_S without a further test, but if and only if the test was negative then he would choose a_R. Hence $t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ (see equation (2.4)).

We have $\overline{c}_{12}^{+}\left(\frac{y}{x}\right) < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ since $Pr(\omega^t) < Pr(\omega^t|+)$.

First period:

If $c > \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then he would neither continue testing nor start testing in the second period, hence

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then $t_1 = Y$: He will test because he would also start testing in the second period.

If $\overline{c}_{12}^{+}\left(\frac{y}{x}\right) < c < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

If $c < \overline{c}_{12}^{+}\left(\frac{y}{x}\right) < \overline{c}_{01}^{S}\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

2.9.2 Lemma 2

Note that the DM's optimal decision rule can be characterised by the set Z_{DM} of unordered histories of test outcomes (see Def A1) such that the DM chooses a_S if and only if $\widetilde{h}_2 \in Z_{DM}$. I divide the range of DM's preferences $\frac{y_{DM}}{x_{DM}}$ into sections such that the optimal decision rule Z_{DM} is constant over any given section, as defined in Def 2a. Note that DM's optimal decision rule is independent of the advisor's testing strategy T. Given Z_{DM}, I solve for T using backward induction.

2.9.2.1 Progressive DM

By definition of progressive,

$$Z_{DM} = \{(++), (+), (+-)\}.$$

Second Period

$t_2^{+} = N$: Given $h_1 = (+)$, the DM will choose a_S without a further test, and if the advisor tested then the DM would choose a_S for either outcome of the second test. Hence, it is optimal not to test.

$t_2^{-} = Y$ if and only if $c < \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$: Given $h_1 = (-)$, the DM will choose a_R without a further test, but if the advisor tests then the DM will choose a_S if and only if the test was positive. Hence, $t_2^{-} = Y$ if and

only if $c < \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ (see equation (2.3)).

$t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$: Given $h_1 = (\emptyset)$, the DM will choose a_R without a further test, but if the advisor tests then the DM will choose a_S if and only if the test was positive. Hence, $t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ (see equation (2.6)).

Clearly, $\overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right) < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$, since $Pr\left(\omega^T|-\right) < Pr\left(\omega^T\right)$.

First period

If $c > \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ then the advisor will neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ the advisor will test because he would also start testing in the second period.

If $c < \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right) < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = Y, t_2^{\emptyset} = Y\right).$$

If $\overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

2.9.2.2 Moderate DM

By definition of moderate,

$$Z_{DM} = \{(++), (+)\}.$$

Second Period

$t_2^{+} = Y$ if and only if $c < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$: Given $h_1 = (+)$, the DM chooses a_S without a further test, but if the advisor tests then the DM would choose a_R if and only if the test was negative. Hence, $t_2^{+} = Y$ if and only if $c < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ (see equation (2.5)).

$t_2^{-} = N$: Given $h_1 = (-)$, the DM chooses a_R without a further test, and if the advisor tests then the DM would choose a_R irrespective of the outcome of the second test. Hence, it is optimal not to test.

$t_2^{\emptyset} = Y$ if and only if $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$: Given $h_1 = (\emptyset)$, the DM chooses a_R without a further test, but if the

advisor tests then the DM would choose a_S if and only if the test was positive. Hence, $t_2^{\emptyset} = Y$ if and only if $c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ (see equation (2.6)).

First period

If $c > max\left\{\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right), \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)\right\}$, then the advisor will neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$ then $t_1 = Y$. The advisor would start testing in the second period, so it must be optimal to test in the first period.

If $c < min\left\{\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right), \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)\right\}$, then

$$T = \left(t_1 = Y, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

If $\overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right)$, then

$$T = \left(t_1 = Y, t_2^{+} = N, t_2^{-} = N, t_2^{\emptyset} = Y\right).$$

If $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$, then without any test the DM will choose a_R but if the advisor tests then the DM chooses a_S if and only if both tests are positive. Hence, $t_1 = Y$ if and only if $c < \overline{c}_{02}^{R}\left(\frac{y_A}{x_A}\right)$ (see equation (2.7)). If $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right) < \overline{c}_{02}^{R}\left(\frac{y_A}{x_A}\right) < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$.

If $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{02}^{R}\left(\frac{y_A}{x_A}\right) < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then

$$T = \left(t_1 = Y, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

If $\overline{c}_{01}^{R}\left(\frac{y_A}{x_A}\right) < \overline{c}_{02}^{R}\left(\frac{y_A}{x_A}\right) < c < \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then

$$T = \left(t_1 = N, t_2^{+} = Y, t_2^{-} = N, t_2^{\emptyset} = N\right).$$

2.9.2.3 Conservative DM

By definition of conservative,

$$Z_{DM} = \{(++)\}.$$

Second Period

$t_2^+ = Y$ if and only if $c < \widetilde{c}^+\left(\frac{y_A}{x_A}\right)$: Given $h_1 = (+)$, the DM chooses a_R without a further test, but if the advisor tests then the DM would choose a_S if and only if the test was positive. Hence, $t_2^+ = Y$ if and only if $c < \widetilde{c}^+\left(\frac{y_A}{x_A}\right)$ (see equation (2.8)).

$t_2^- = N$: Given $h_1 = (-)$, the DM chooses a_R without a further test, and if the advisor tests then the DM would choose a_R irrespective of the outcome of the second test. Hence, it is optimal not to test.

$t_2^{\emptyset} = N$: Given $h_1 = (\emptyset)$, the DM chooses a_R without a further test, and if the advisor tests then the DM would choose a_R irrespective of the outcome of the test. Hence, it is optimal not to test.

First period

If $c > \widetilde{c}^+\left(\frac{y_A}{x_A}\right)$ then the advisor will neither continue testing nor start testing in the second period, hence it cannot be optimal to test in the first period:

$$T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right).$$

If $c < \widetilde{c}^+\left(\frac{y_A}{x_A}\right)$ then $t_1 = Y$. Without a test the DM chooses a_R, but if the advisor tests then the DM will choose a_S if and only if both tests are positive. Hence, $t_1 = Y$ if and only if $c < \overline{c}_{02}^R\left(\frac{y_A}{x_A}\right)$ (see equation (2.7)). Clearly, $\overline{c}_{02}^R\left(\frac{y}{x}\right) < \widetilde{c}^+\left(\frac{y}{x}\right)$.

If $c < \overline{c}_{02}^R\left(\frac{y}{x}\right) < \widetilde{c}^+\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = Y, t_2^+ = Y, t_2^- = N, t_2^{\emptyset} = N\right).$$

If $\overline{c}_{02}^R\left(\frac{y}{x}\right) < c < \widetilde{c}^+\left(\frac{y}{x}\right)$ then

$$T = \left(t_1 = N, t_2^+ = Y, t_2^- = N, t_2^{\emptyset} = N\right).$$

2.9.3 Lemma 3

I will again divide the range of the DM's preferences into sections, labelled as in Def 2a and 2b. To summarize each player's preferences, I will write a hypothetical decision rule for each player characterising how he would behave as a single player under observable testing. In particular, let the set Z_i of unordered histories of test outcomes (see Def A1) be such that player i would choose a_S if and only if $\widetilde{h}_2 \in Z_i$ assuming he can observe $\widetilde{h}_2$.

Define a belief assessment b_{FD} such that the DM's belief are as they would if the advisor's message fully discloses outcomes, i.e. if $m = \widetilde{h}_2$ for any h_2.

Definition A2

Define belief assessment b_{FD} as follows:

$$\begin{aligned} Pr(h_2 = (\emptyset\emptyset) \,|m = (\emptyset)) &= 1 \\ Pr(h_2 = (++) \,|m = (++)) &= 1 \\ Pr(h_2 = (--) \,|m = (--)) &= 1 \\ Pr(h_2 = (+\emptyset) \,|m = (+)) &= 1 \\ Pr(h_2 = (-\emptyset) \,|m = (-)) &= 1 \end{aligned}$$

If $\frac{y_A}{x_A} > \frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))}$ *then*

$$Pr(h_2 = (+-) \,|m = (+-)) = 1$$

If $\frac{y_A}{x_A} \leq \frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))}$ *then*

$$Pr(h_2 = (-+) \,|m = (+-)) = 1.$$

Claim 1: If the DM has belief assessment b_{FD} then her optimal decision rule is such that she chooses a_S if and only if $m = \widetilde{h}_2 \in Z_{DM}$, i.e. $\Sigma_{DM} = Z_{DM}$.

Proof Claim 1: If the DM has belief assessment b_{FD} then given message $m = \widetilde{h}_2'$ she assigns the same posterior belief to the event that $\omega = \omega^i$ and $h_2 = h_2'$ for each state $i = T, F$ and any h_2 as she does if she observed $h_2 = h_2'$, for any given h_2'. Hence, the set of unordered histories of outcomes for which she optimally chooses a_S under observable setting is equal to the set of messages for which she optimally chooses a_S.

Claim 2: Suppose the DM's beliefs are $b = b_{FD}$ and her decision rule is characterised by $\Sigma_{DM} = Z_{DM}$. If the advisor's reporting strategy is such that $m = \widetilde{h}_2$ for any given h_2 then there exists an equilibrium in

which he uses the same testing strategy as under observable testing.

Proof Claim 2:

Given $b = b_{FD}$, it is optimal for the DM to choose a_S if and only if $m \in \Sigma_{DM} = Z_{DM}$ by Claim 1.

Given the DM's strategy, if the advisor's strategy were to report $m = \widetilde{h}_2$ for any h_2 then the DM chooses a_S if and only if she would have chosen a_S conditional on observing h_2. Thus, if the advisor's strategy were to report $m = \widetilde{h}_2$ for any h_2 then his testing strategy under observable testing is a best response to the DM's strategy.

Finally, if the advisor's strategy were to report $m = \widetilde{h}_2$ for any h_2 then the belief assessment b_{FD} is consistent with the testing and reporting strategy. Given this reporting strategy, $m = (--)$ only if $h_2 = (--)$, $m = (++)$ only if $h_2 = (++)$ and $m = (\emptyset)$ only if $h_2 = (\emptyset\emptyset)$. In addition, given that there is no optimal testing strategy under observable testing such that the advisor only tests in the second period, $m = (+)$ only if $h_2 = (+\emptyset)$ and $m = (-)$ only if $h_2 = (-\emptyset)$. And lastly, an advisor with $\frac{y_A}{x_A} < \frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))}$ only tests in the second period following a negative outcome hence $m = (+-)$ only if $h_2 = (-+)$, while an advisor with $\frac{y_A}{x_A} \geq \frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))}$ only tests in the second period following a positive outcome hence $m = (+-)$ only if $h_2 = (+-)$.

Claim 3: If $Z_A = Z_{DM}$, i.e. if the advisor is unbiased, then the following is an equilibrium: The DM holds beliefs $b = b_{FD}$ and follows a decision rule characterised by $\Sigma_{DM} = Z_{DM}$. The advisor reports $m = \widetilde{h}_2$ for any h_2 and follows the same testing strategy as under observable testing (see proof of Lemma 2 in Section 2.9.2).

Proof Claim 3: Given Claim 1 and 2, it is sufficient to show that the advisor reports $m = \widetilde{h}_2$ for any h_2. The advisor will optimally send $m = \widetilde{h}_2$ or any h_2 because $\widetilde{h}_2 \in Z_A$ if and only if $\widetilde{h}_2 \in \Sigma_{DM}$ given $Z_A = Z_{DM} = \Sigma_{DM}$.

2.9.3.1 Progressive DM

Assume $p_1 > p_2$. If the DM is progressive then $Z_{DM} = \{(+-), (+), (++)\}$.

2.9.3.1.1) Advisor such that $Z_A \supset Z_{DM} = \{(+-),(+),(++)\}$. $Z_A \supset Z_{DM}$ implies that the advisor is more enthusiastic than the DM. $Z_A \supset Z_{DM}$ if and only if the advisor's preferences satisfy

$$\frac{y_A}{x_A} \leq \frac{q}{1-q}.$$

.Equilibrium: The DM holds beliefs $b = b_{FD}$ and follows a decision rule characterised by $\Sigma_{DM} = Z_{DM}$. The advisor reports $m = \widetilde{h}_2$ for any h_2 and follows the same testing strategy as under observable testing.

Given Claim 1 and 2, it is sufficient to show that the advisor reports $m = \widetilde{h}_2$ for any h_2. If h_2 is such that $\widetilde{h}_2 \in Z_{DM}$ then the advisor will optimally send $m = \widetilde{h}_2$ because $Z_A \supset Z_{DM}$. In addition, if h_2 is such that $\widetilde{h}_2 \notin Z_{DM}$, i.e. $\widetilde{h}_2 \in \{(--),(-),(\emptyset)\}$, there is no feasible m such that $m \in Z_{DM}$. Therefore, the advisor weakly prefers $m = \widetilde{h}_2$ for any h_2.

2.9.3.1.2) Advisor such that $Z_A = Z_{DM} = \{(+-),(+),(++)\}$. $Z_A = Z_{DM}$ implies that the advisor is unbiased. $Z_A = Z_{DM}$ if and only if the advisor's preferences satisfy

$$\frac{q}{1-q} \leq \frac{y_A}{x_A} < \frac{q\left(p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(p_2\left(1-p_2\right)\right)}.$$

See Claim 3.

2.9.3.1.3) Advisor such that $Z_A = \{(++),(+)\} \subset Z_{DM} = \{(+-),(+),(++)\}$. $Z_A = \{(++),(+)\}$ implies that the advisor is more reluctant. $Z_A = \{(++),(+)\}$ if and only if the advisor's preferences satisfy

$$\frac{q\left(p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(p_2\left(1-p_2\right)\right)} \leq \frac{y_A}{x_A} < \frac{qp_1}{\left(1-q\right)\left(1-p_2\right)}.$$

Equilibrium:

- The DM's belief assessment $b = b_1$ is as follows:

  If $c \leq \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right)$,

$$Pr\left(h_2 = (-\emptyset) \,|m = (-)\right) = \frac{Pr\left(h_2 = (-\emptyset)\right)}{Pr\left(h_2 = (-\emptyset)\right) + Pr\left(h_2 = (+-)\right)}$$

$$Pr\left(h_2 = (+-) \,|m = (-)\right) = \frac{Pr\left(h_2 = (+-)\right)}{Pr\left(h_2 = (-\emptyset)\right) + Pr\left(h_2 = (+-)\right)}.$$

If $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$,

$$Pr\left(h_2 = (-\emptyset)\,|m = (-)\right) = 1.$$

Given any $m \neq (-)$ her assessment is as in $b = b_{FD}$ (see Def A2).

- the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$,
- the advisor's testing strategy is as in the case of a single player (see Section 2.9.1.4),
- the advisor's reporting strategy specifies $m = (-)$ for either $h_2 = (+-)$ or $h_2 = (-+)$ and $m = \widetilde{h}_2$ otherwise.

The DM's decision rule is optimal: Given beliefs b_1, the DM optimally chooses a_R given $m = (-)$ since if $c \leq \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(-)\right)}{Pr\left(\omega^F|m=(-)\right)} &= \frac{Pr\left(h_2=(+-)\,|\omega^T\right)+Pr\left(h_2=(-\emptyset)\,|\omega^T\right)}{Pr\left(h_2=(+-)\,|\omega^F\right)+Pr\left(h_2=(-\emptyset)\,|\omega^F\right)} \\ &= \frac{q\left(1-p_1+p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(p_2+p_2\left(1-p_2\right)\right)} < \frac{q}{1-q} < \frac{y_{DM}}{x_{DM}},\end{aligned}$$

and if $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(-)\right)}{Pr\left(\omega^F|m=(-)\right)} &= \frac{Pr\left(h_2=(-\emptyset)\,|\omega^T\right)}{Pr\left(h_2=(-\emptyset)\,|\omega^F\right)} \\ &= \frac{q\left(1-p_1\right)}{\left(1-q\right)p_2} < \frac{q}{1-q} < \frac{y_{DM}}{x_{DM}}.\end{aligned}$$

In addition, if $m \neq (-)$ then she optimally chooses a_S if and only if $m \in Z_{DM}$, since for $m \neq (-)$ her beliefs are as for $b = b_{FD}$ and Claim 1 applies. Hence, for all m the DM optimally chooses a_S if and only if $m \in Z_{DM}$. Hence, $\Sigma_{DM} = Z_{DM}$.

The advisor's reporting strategy is optimal given the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$. If he sends $m = \widetilde{h}_2$ then the DM acts in his interest for any h_2 provided h_2 is such that $\widetilde{h}_2 \neq (+-)$. If h_2 is such that $\widetilde{h}_2 = (+-)$ then by sending $m = (-)$ he causes the DM to choose a_R and $\widetilde{h}_2 = (+-) \notin Z_A$.

The advisor's optimal testing strategy in the case of a single player is also optimal in this case: Given that $\Sigma_{DM} = Z_{DM}$ and given the advisor's reporting strategy, for any h_2 the DM chooses a_S if and only if $\widetilde{h}_2 \in Z_A$, i.e. the DM behaves as the advisor would behave for any h_2.

The DM's beliefs are consistent because if $c \leq \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then $m = (-)$ only if either $h_2 = (-\emptyset)$ or

$h_2 = (+-)$, and if $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then $m = (-)$ only if $h_2 = (-\emptyset)$. For any c, $m = (+)$ only if $h_2 = (+\emptyset)$, $m = (++)$ only if $h_2 = (++)$, $m = (\emptyset)$ only if $h_2 = (\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

2.9.3.1.4) Advisor such that $Z_A = \{(++)\} \subset Z_{DM} = \{(+-),(+),(++)\}$. $Z_A = \{(++)\}$ implies that the advisor is more reluctant. $Z_A = \{(++)\}$ holds if and only if the advisor's preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)} \leq \frac{y_A}{x_A} < \frac{qp_1}{(1-q)(1-p_2)^2}.$$

Equilibrium:

- The DM's belief assessment is $b = b_1$ as defined above (see Section A.3.1.3),
- the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$,
- the advisor's testing strategy is as in the case of a single player (see Section 2.9.1.5),
- the advisor's reporting strategy specifies $m = (-)$ for either $h_2 = (+-)$ or $h_2 = (-+)$, and $m = (\emptyset)$ for either $h_2 = (+\emptyset)$ or $h_2 = (\emptyset+)$, and $m = \widetilde{h}_2$ otherwise.

The DM's decision rule characterised by $\Sigma_{DM} = Z_{DM}$ is optimal given $b = b_1$ for the same argument as in Section A.3.1.3.

The advisor's reporting strategy is optimal given the DM's strategy $\Sigma_{DM} = Z_{DM}$: If he sends $m = \widetilde{h}_2$ then the DM acts in his interest for any h_2 provided h_2 is such that $\widetilde{h}_2 \neq (+-)$ and $\widetilde{h}_2 \neq (+)$. If h_2 is such that $\widetilde{h}_2 = (+-)$ then by sending $m = (-)$ he causes the DM to choose a_R and $\widetilde{h}_2 = (+-) \notin Z_A$. If h_2 is such that $\widetilde{h}_2 = (+)$ then by sending $m = (\emptyset)$ he causes the DM to choose a_R and $\widetilde{h}_2 = (+) \notin Z_A$.

The advisor's optimal strategy in the case of a single player is also optimal in this case: Given that $\Sigma_{DM} = Z_{DM}$ and given the advisor's reporting strategy, for any h_2 the DM chooses a_S if and only if $\widetilde{h}_2 \in Z_A$, i.e. the DM behaves as the advisor would behave for any h_2.

The DM's beliefs are consistent because $m = (-)$ only if either $h_2 = (-\emptyset)$ or $h_2 = (+-)$, $m = (++)$ only if $h_2 = (++)$, $m = (\emptyset)$ only if $h_2 = (\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

2.9.3.1.5) Advisor such that Z_A is empty. If Z_A is empty then this implies that the advisor is more reluctant. Z_A is empty if and only if the advisor's preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)^2} \leq \frac{y_A}{x_A}.$$

Equilibrium:

- the DM's belief assessment is $b = b_{FD}$,
- the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$,
- and the advisor's testing strategy is $T = \left(t_1 = N, t_2^+ = N, t_2^- = N, t_2^{\emptyset} = N\right)$,
- the advisor sends $m = (\emptyset)$ for any h_2.

The DM's decision rule is optimal given beliefs $b = b_{FS}$ by Claim 1.

The advisor's reporting strategy is optimal given the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$: Sending $m = (\emptyset)$ for any h_2 is optimal for the advisor because the DM chooses a_R for any h_2, which is what the advisor prefers.

The advisor's testing strategy is optimal: he always achieves his preferred action a_R by selecting not to test since $\Sigma_{DM} = Z_{DM}$ and $m = (\emptyset)$ for any h_2.

The DM's belief assessment is consistent because $m = (\emptyset)$ only if $h_2 = (\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

2.9.3.2 Moderate DM

If the DM is moderate then $Z_{DM} = \{(+), (++)\}$.

2.9.3.2.1) Advisor such that $Z_A \supset Z_{DM} = \{(+), (++)\}$. $Z_A \supset Z_{DM}$ implies that the advisor is more enthusiastic. $Z_A \supset Z_{DM}$ holds if and only if

$$\frac{y_A}{x_A} \leq max\left\{\frac{q(p_1(1-p_1))}{(1-q)(p_2(1-p_2))}, \frac{q}{1-q}\right\}.$$

In this case, I distinguish between moderate leaning towards progressive and moderate leaning towards conservative.

If the DM is moderate leaning towards progressive:

Equilibrium:

- the DM's belief assessment $b = b_2$ is as follows:

  If $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$,

$$Pr(h_2 = (+\emptyset)\,|m = (+)) = \frac{Pr(h_2 = (+\emptyset))}{Pr(h_2 = (+\emptyset)) + Pr(h_2 = (-+))}$$

$$Pr(h_2 = (-+)\,|m = (+)) = \frac{Pr(h_2 = (-+))}{Pr(h_2 = (+\emptyset)) + Pr(h_2 = (-+))}.$$

  If $c > \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$,

$$Pr(h_2 = (+\emptyset)\,|m = (+)) = 1.$$

  Given any $m \neq (+)$ her assessment is as in $b = b_{FD}$ (see Def A2).

- The DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$,

- the advisor's testing strategy is as in the case of observable testing when the DM is progressive (see Section 2.9.2.1),

- the advisor sends $m = (+)$ if $h_2 = (-+)$ or $h_2 = (+-)$ and otherwise he sends $m = \widetilde{h}_2$ for any h_2.

The DM's decision rule is optimal: Given beliefs b_1, by definition of "leaning towards conservative" the DM optimally chooses a_S given $m = (+)$ since if $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(+)\right)}{Pr\left(\omega^F|m=(+)\right)} &= \frac{Pr\left(h_2=(-+)\,|\omega^T\right)+Pr\left(h_2=(+\emptyset)\,|\omega^T\right)}{Pr\left(h_2=(-+)\,|\omega^F\right)+Pr\left(h_2=(+\emptyset)\,|\omega^F\right)} \\ &= \frac{q\left(p_1+p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(1-p_2+p_2\left(1-p_2\right)\right)} > \frac{y_{DM}}{x_{DM}},\end{aligned}$$

and if $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(+)\right)}{Pr\left(\omega^F|m=(+)\right)} &= \frac{Pr\left(h_2=(+\emptyset)\,|\omega^T\right)}{Pr\left(h_2=(+\emptyset)\,|\omega^F\right)} \\ &= \frac{qp_1}{\left(1-q\right)\left(1-p_2\right)} > \frac{y_{DM}}{x_{DM}}.\end{aligned}$$

If $m \neq (+)$ then she optimally chooses a_S if and only if $m \in Z_{DM}$, since for $m \neq (+)$ her beliefs are as

for $b = b_{FD}$ and Claim 1 applies. Hence, for all m the DM optimally chooses a_S if and only if $m \in Z_{DM}$. Hence, $\Sigma_{DM} = Z_{DM}$.

The advisor's reporting strategy is optimal given the DM's strategy $\Sigma_{DM} = Z_{DM}$: If h_2 is such that $\widetilde{h}_2 \in Z_{DM}$ the advisor will optimally send $m = \widetilde{h}_2$ because $Z_A \supset Z_{DM}$. If h_2 is such that $\widetilde{h}_2 = (-+)$ then by sending $m = (+)$ he causes the DM to choose a_S and $(-+) \in Z_A$. In addition, if h_2 is such that $\widetilde{h}_2 \notin Z_{DM}$ and $\widetilde{h}_2 \neq (-+)$ there is no feasible m such that $m \in Z_{DM}$.

The advisor's optimal strategy when the DM is progressive and testing is observable is also a best response in this case. Given that $\Sigma_{DM} = Z_{DM}$ and given the advisor's reporting strategy, the advisor achieves a_S if and only if $h_2 \in \{(-+),(+-),(\emptyset+),(+\emptyset),(++)\}$. This is the same as if h_2 was observable and the DM was progressive.

The DM's beliefs are consistent because if $c \leq \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$ then $m = (+)$ only if either $h_2 = (+\emptyset)$ or $h_2 = (-+)$, and if $c > \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$ then $m = (+)$ only if $h_2 = (+\emptyset)$. For any c, $m = (-)$ only if $h_2 = (-\emptyset)$, $m = (--)$ only if $h_2 = (--)$, $m = (\emptyset)$ only if $h_2 = (\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

If the DM is moderate leaning towards conservative:

Equilibrium:

- the DM's belief assessment $b = b_2$, as defined above for a moderate leaning progressive DM,
- The DM's strategy is characterised by $\Sigma_{DM} = \{(++)\} \subset Z_{DM}$ if $c \leq \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$ and by $\Sigma_{DM} = Z_{DM}$ if $c > \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$,
- the advisor's testing strategy is as in the case of observable testing when the DM is progressive if $c > \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$, and as in the case of observable testing when the DM is conservative if $c \leq \overline{c}^{-}_{12}\left(\frac{y_A}{x_A}\right)$ (see Section 2.9.2.1),
- the advisor sends $m = (+)$ if either $h_2 = (-+)$ or $h_2 = (+-)$ and otherwise he sends $m = \widetilde{h}_2$ for any h_2.

The DM's decision rule is optimal: Given beliefs b_1, by definition of "leaning towards conservative" the

DM optimally chooses a_R given $m=(+)$ since if $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(+)\right)}{Pr\left(\omega^F|m=(+)\right)} &= \frac{Pr\left(h_2=(-+)|\omega^T\right)+Pr\left(h_2=(+\emptyset)|\omega^T\right)}{Pr\left(h_2=(-+)|\omega^F\right)+Pr\left(h_2=(+\emptyset)|\omega^F\right)} \\ &= \frac{q\left(p_1+p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(1-p_2+p_2\left(1-p_2\right)\right)} < \frac{y_{DM}}{x_{DM}},\end{aligned}$$

and optimally chooses a_S given $m=(+)$ if $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$

$$\begin{aligned}\frac{Pr\left(\omega^T|m=(+)\right)}{Pr\left(\omega^F|m=(+)\right)} &= \frac{Pr\left(h_2=(+\emptyset)|\omega^T\right)}{Pr\left(h_2=(+\emptyset)|\omega^F\right)} \\ &= \frac{qp_1}{\left(1-q\right)\left(1-p_2\right)} > \frac{y_{DM}}{x_{DM}}.\end{aligned}$$

If $m \neq (+)$ then she optimally chooses a_S if and only if $m \in Z_{DM}$, since for $m \neq (+)$ her beliefs are as for $b = b_{FD}$ and Claim 1 applies. Hence, if $c > \overline{c}_{12}^{+}\left(\frac{y_A}{x_A}\right)$ then for all m the DM optimally chooses a_S if and only if $m \in Z_{DM}$ and hence, $\Sigma_{DM} = Z_{DM}$. But if $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ then $\Sigma_{DM} = \{(++)\} \subset Z_{DM}$.

The advisor's reporting strategy is optimal if $c > \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ given the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$: If h_2 is such that $\widetilde{h}_2 \in Z_{DM}$ then the advisor will optimally send $m = \widetilde{h}_2$ because $Z_A \supset Z_{DM}$. If h_2 is such that $\widetilde{h}_2 = (-+)$ then by sending $m=(+)$ he causes the DM to choose a_S and $\widetilde{h}_2 = (-+) \in Z_A$. In addition, if $\widetilde{h}_2 \notin Z_{DM}$ and $\widetilde{h}_2 \neq (-+)$ there is no feasible m such that $m \in Z_{DM}$.

The advisor's reporting strategy is also optimal if $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ given the DM's strategy is characterised by $\Sigma_{DM} = \{(++)\}$. If $h_2 = (++)$ the advisor will optimally send $m = \widetilde{h}_2$ because $(++) \in Z_A \supset Z_{DM}$. In addition, if h_2 is such that $\widetilde{h}_2 \notin Z_{DM}$ there is no feasible m such that $m \in Z_{DM}$. If $h_2 = (+-)$ the advisor is indifferent between sending $m = \widetilde{h}_2$ and sending $m=(+)$ because either triggers a_R.

The DM's beliefs are consistent because if $c \leq \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ then $m=(++)$ only if $h_2=(++)$, $m=(-)$ only if $h_2=(-\emptyset)$, $m=(\emptyset)$ only if $h_2=(\emptyset\emptyset)$ and no other messages are sent on the equilibrium path. And if $c > \overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right)$ then $m=(+)$ only if $h_2=(+\emptyset)$, $m=(-)$ only if $h_2=(-\emptyset)$, $m=(\emptyset)$ only if $h_2=(\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

2.9.3.2.2) Advisor such that $Z_A = Z_{DM} = \{(+),(++)\}$. $Z_A = Z_{DM}$ implies that the advisor is unbiased. $Z_A = Z_{DM}$ holds if and only if

$$max\left\{\frac{q\left(p_1\left(1-p_1\right)\right)}{\left(1-q\right)\left(p_2\left(1-p_2\right)\right)}, \frac{q}{1-q}\right\} \leq \frac{y_A}{x_A} < \frac{qp_1}{\left(1-q\right)\left(1-p_2\right)}.$$

See Claim 3.

2.9.3.2.3) Advisor such that $Z_A = \{(++)\} \subset Z_{DM} = \{(+),(++)\}$. $Z_A = \{(++)\}$ implies that the advisor is more reluctant and holds if and only if

$$\frac{qp_1}{(1-q)(1-p_2)} \leq \frac{y_A}{x_A} < \frac{qp_1}{(1-q)(1-p_2)^2}.$$

Equilibrium:

- The DM's belief assessment is $b = b_{FD}$,
- the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$,
- and the advisor's testing strategy is as in the case of a single player (see Section 2.9.1.5),
- the advisor sends $m = (\emptyset)$ if $h_2 = (+\emptyset)$ or $h_2 = (\emptyset+)$ and otherwise sends $m = \widetilde{h}_2$.

The DM's decision rule is optimal given beliefs $b = b_{FS}$ by Claim 1.

The advisor's reporting strategy is optimal given the DM's strategy is characterised by $\Sigma_{DM} = Z_{DM}$: If he sends $m = \widetilde{h}_2$ then the DM acts in his interest for any h_2 provided h_2 is such that $\widetilde{h}_2 \neq (+)$. If h_2 is such that $\widetilde{h}_2 = (+)$ then by sending $m = (\emptyset)$ he causes the DM to choose a_R and $\widetilde{h}_2 = (+) \notin Z_A$.

The advisor's optimal strategy in the case of a single player is also optimal in this case: Given that $\Sigma_{DM} = Z_{DM}$ and given the advisor's reporting strategy, for any $\widetilde{h}_2$ the DM chooses a_S if and only if $\widetilde{h}_2 \in Z_A$, i.e. the DM behaves as the advisor would behave for any h_2.

The DM's beliefs are consistent because $m = (-)$ only if $h_2 = (-\emptyset)$, $m = (++)$ if and only if $h_2 = (++)$ and $m = (\emptyset)$ if $h_2 = (\emptyset\emptyset)$ and no other messages are sent on the equilibrium path.

2.9.3.2.4) Advisor such that Z_A is empty. If Z_A is empty then this implies that the advisor is more reluctant. Z_A is empty if and only if the advisor's preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)^2} \leq \frac{y_A}{x_A}.$$

The equilibrium analysed in Section 2.9.3.1.5) is also an equilibrium here because the argument is independent of the DM's type.

2.9.3.3 Conservative DM

If the DM is conservative then $Z_{DM} = \{(++)\}$.

A.2.9.3.3.1) Advisor such that $Z_A \subset Z_{DM} = \{(++)\}$. $Z_A \subset Z_{DM}$ implies that the advisor is more enthusiastic. $Z_A \subset Z_{DM}$ holds if and only if the advisor's preferences satisfy

$$\frac{y_A}{x_A} \leq \frac{qp_1}{(1-q)(1-p_2)}.$$

Equilibrium: The DM holds beliefs $b = b_{FD}$ and follows a decision rule characterised by $\Sigma_{DM} = Z_{DM}$. The advisor reports $m = \widetilde{h}_2$ for any h_2 and follows the same testing strategy as under observable testing.

Given Claim 1 and 2, it is sufficient to show that the advisor reports $m = \widetilde{h}_2$ for any $\widetilde{h}_2$. If $h_2 = (++)$ then the advisor will optimally send $m = \widetilde{h}_2$ because the DM chooses a_S and $(++) \in Z_A$. Otherwise, there exists no feasible message such that $m \in Z_{DM}$. Therefore, the advisor weakly prefers $m = \widetilde{h}_2$ for any h_2.

A.2.9.3.3.2) Advisor such that $Z_A = Z_{DM} = \{(++)\}$. $Z_A = Z_{DM}$ implies the advisor is unbiased and holds if and only if the advisor's preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)} \leq \frac{y_A}{x_A} < \frac{qp_1}{(1-q)(1-p_2)^2}.$$

See Claim 3.

A.2.9.3.3.3) Advisor such that Z_A is empty. If Z_A is empty then this implies that the advisor is more reluctant. Z_A is empty if and only if the advisor's preferences satisfy

$$\frac{qp_1}{(1-q)(1-p_2)^2} < \frac{y_A}{x_A}.$$

The equilibrium analysed in Section 2.9.3.1.5) is also an equilibrium here because the argument is independent of the DM's type.

2.9.4 Theorem 1

The following remark ranks the DM's expected loss L_{DM} for different behaviours of the DM and the advisor in equilibrium.

Remark A1

Given an progressive DM,

$$L_{DM}\left(T2^{-}_{(+),(-+)}\right) < L_{DM}\left(T1_{(+)}\right) < L_{DM}\left(T2^{+}_{(++)}\right) < L_{DM}\left(T0\right)$$

Given a moderate DM,

$$L_{DM}\left(T2^{+}_{(++)}\right) < L_{DM}\left(T1_{(+)}\right) < L_{DM}\left(T0\right) < L_{DM}\left(T2^{-}_{(+),(-+)}\right)$$

Given a conservative DM,

$$L_{DM}\left(T2^{+}_{(++)}\right) < L_{DM}\left(T0\right) < L_{DM}\left(T1_{(+)}\right) < L_{DM}\left(T2^{-}_{(+),(-+)}\right)$$

Proof Remark A1: To rank the DM's loss for any equilibrium play we can compare her decision rule to the one she would implement if she saw two test outcomes. The DM's loss increases with the number of possible realisations of h_2 for which she deviates from her optimal decision rule.

If the DM is progressive, $T2^{-}_{(+),(-+)}$ allows her to choose optimally for any h_2. If $T1_{(+)}$ then she deviates from the optimal decision rule if and only if $h_2 = (-+)$. If $T2^{+}_{(++)}$ then she deviates from the optimal decision rule if and only if $h_2 \in \{(-+),(+-)\}$. Finally, if $T0$ she deviates from the optimal decision rule if and only if $h_2 \in \{(-+),(+-),(++)\}$.

If the DM is moderate, $T2^{+}_{(++)}$ allows her to choose optimally for any h_2. If $T1_{(+)}$ then she deviates from the optimal decision rule if and only if $h_2 = (+-)$. If $T0$ then she deviates from the from the optimal decision rule if and only if $h_2 = (++)$. However, $L_{DM}\left(T1_{(+)}\right) < L_{DM}\left(T0\right)$ if and only if $\frac{y}{x} \leq \frac{qp_2}{(1-q)(1-p_1)}$ since $0 < \overline{c}^{R}_{01}\left(\frac{y}{x}\right)$ for $\frac{y}{x} \leq \frac{qp_2}{(1-q)(1-p_1)}$ (see Section 2.9.1.4).If $T2^{-}_{(+),(+-)}$then she deviates from the optimal decision rule if and only if $h_2 \in \{(-+),(+-)\}$.

If the DM is conservative, $T2^{+}_{(++)}$ allows her to choose optimally for any h_2. If $T0$ she deviates from

the optimal decision rule if and only if $h_2 = (++)$. If $T1_{(+)}$ then she deviates from the optimal decision rule if and only if $h_2 = (+-)$. However, $L_{DM}(T0) < L_{DM}\left(T1_{(+)}\right)$ if and only if $\frac{y}{x} > \frac{qp_2}{(1-q)(1-p_1)}$ since $0 > \overline{c}^R_{01}\left(\frac{y}{x}\right)$ for $\frac{y}{x} > \frac{qp_2}{(1-q)(1-p_1)}$ (see Section 2.9.1.4). If $T2^-_{(+),(+-)}$ then she deviates from the optimal decision rule if and only if $h_2 \in \{(-+),(+-)\}$.

2.9.4.1 Progressive DM

If $\overline{c}^R_{01}\left(\frac{y_A}{x_A}\right) < c < \overline{c}^R_{02}\left(\frac{y_A}{x_A}\right)$, then behaviour is characterised by $T0$ under observable testing and by $T2^+_{(++)}$ under hidden testing by Lemma 1 and 2. By Remark A1, the DM is better off.

If $c < min\left\{\overline{c}^+_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^R_{01}\left(\frac{y_A}{x_A}\right)\right\}$, then behaviour is characterised by $T1_{(+)}$ under observable testing and by $T2^+_{(++)}$ under hidden testing by Lemma 1 and 2. By Remark A1, the DM is worse off.

Otherwise, behaviour in equilibrium is the same irrespective of whether testing is observable or hidden.

2.9.4.2 Moderate leaning towards progressive DM

If $c < \overline{c}^-_{12}\left(\frac{y_A}{x_A}\right)$ behaviour is characterised by $T1_{(+)}$ under observable testing vs $T2^-_{(+),(-+)}$ under hidden testing by Lemma 1 and 2. By Remark A1, the DM is worse off.

Otherwise, behaviour in equilibrium is the same irrespective of whether testing is observable or hidden.

2.9.4.3 Moderate leaning towards conservative DM

If $c < min\left\{\overline{c}^-_{12}\left(\frac{y_A}{x_A}\right), \overline{c}^R_{02}\left(\frac{y_A}{x_A}\right)\right\}$ then behaviour is characterised by $T1_{(+)}$ under observable testing and by $T2^+_{(++)}$ under hidden testing by Lemma 1 and 2. By Remark A1, the DM is better off.

If $\overline{c}^R_{02}\left(\frac{y_A}{x_A}\right) < c < \overline{c}^-_{12}\left(\frac{y_A}{x_A}\right)$ then behaviour is characterised by $T1_{(+)}$ under observable testing and by $T0$ under hidden testing by Lemma 1 and 2. By Remark A1, the DM is worse off.

Otherwise, behaviour in equilibrium is the same irrespective of whether testing is observable or hidden.

2.9.4.4 Conservative DM

By Lemma 1 and 2, behaviour in equilibrium is the same irrespective of whether testing is observable or hidden.

2.9.5 Proposition 1

2.9.5.1 Progressive DM

Suppose $c < min\left\{\overline{c}_{02}^{R}, \overline{c}_{01}^{R}\right\}$. If the DM commits to choose a_R if $\widetilde{h}_2 = (+-)$ then her decision rule is characterised by the set $Z_{DM} = \{(++),(+)\}$ of unordered histories of test outcomes. Hence, the DM behaves as if she was moderate. Behaviour is characterised by $T2^{+}_{(++)}$ either if testing is observable and the DM is moderate (Lemma 2) or if testing is hidden and the DM is progressive (Lemma 3). Hence, with commitment she achieves the same payoff as under hidden testing. Suppose the DM delegates the decision to the advisor then the advisor's behaviour is also characterised by $T2^{+}_{(++)}$ (Lemma 1).

2.9.5.2 Moderate leaning towards conservative DM

Suppose $c < min\left\{\overline{c}_{12}^{-}\left(\frac{y_A}{x_A}\right), \overline{c}_{02}^{R}\left(\frac{y_A}{x_A}\right)\right\}$. If the DM commits to choose a_R if $\widetilde{h}_2 = (+)$ then her optimal decision rule is characterised by $Z_{DM} = \{(++)\}$. Hence, the DM behaves as if she was conservative. Behaviour is characterised by $T2^{+}_{(++)}$ either if testing is observable and the DM is conservative (Lemma 2) or if testing is hidden and the DM is moderate leaning towards conservative (Lemma 3). Hence, with commitment she achieves the same payoff as under hidden testing. Suppose the DM delegates the decision to the advisor then the advisor's behaviour is never characterised by $T2^{+}_{(++)}$ (Lemma 1) and the DM is worse off by Remark A1.

2.9.6 Proposition 2

By Lemma 2 and 3, if the DM is moderate leaning towards conservative and $\overline{c}_{02}^{R} < c < \overline{c}_{12}^{-}$, then equilibrium play is $T1_{(+)}$ under observable testing and $T0$ under hidden testing. By Remark A1, the DM is strictly worse off. The agent is also strictly worse off by a revealed preference argument: he could have achieved $T0$ by not testing in either period in a game with observable testing, yet by proof of Lemma 2 he strictly prefers to follow his optimal strategy, and therefore he must strictly prefer $T1_{(+)}$.

2.9.7 Proposition 3

2.9.7.1 Progressive DM

By Lemma 2, as y_A decreases from ∞ to 0, players' behaviour in equilibrium either I) remains at $T0$, II) changes from $T0$ to $T1_{(+)}$ or III) changes from $T0$ to $T1_{(+)}$ and then to $T2^{-}_{(+),(-+)}$ depending on the

level of cost. The result follows by Remark A1.

2.9.7.2 Moderate DM

Suppose $c > \widehat{c}$, where $\widehat{c} = \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$. By Lemma 2, as y_A decreases from ∞ to 0, players' behaviour in equilibrium either I) remains at $T0$ or II) changes from $T0$ to $T1_{(+)}$ depending on the level of cost. The result follows by Remark A1.

Suppose $c \leq \widehat{c}$. By Lemma 2, as y_A decreases from ∞ to 0, players' behaviour in equilibrium changes from $T0$ to $T2^{+}_{(++)}$ and then to $T1_{(+)}$. By Remark A1, the DM's preferences over y_A are single-peaked. Further, there exists some c such that the players' behaviour is characterised by $T1_{(+)}$ when the advisor is unbiased and characterised by $T2^{+}_{(++)}$ when the advisor is more reluctant. This is because there are levels of c which satisfy $c > \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right)$ if $\frac{qp_1}{(1-q)(1-p_2)} < \frac{y_A}{x_A} < \frac{qp_1}{(1-q)(1-p_2)^2}$ and also satisfy $c < min\left\{\overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right), \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right)\right\}$ if $\frac{y_A}{x_A} > \frac{qp_1}{(1-q)(1-p_2)^2}$. Hence, by Remark A1, there exists some c such that the DM is strictly better off with a more reluctant advisor than with an unbiased advisor.

2.9.7.3 Conservative DM

By Lemma 2, as y_A decreases from ∞ to 0, players' behaviour in equilibrium either I) remains at $T0$ or II) changes from $T0$ to $T2^{+}_{(++)}$ depending on the level of cost. The result follows by Remark A1.

2.9.8 Proposition 4

2.9.8.1 Progressive DM

Suppose $c > \widehat{c}$, where $\widehat{c} = \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$. By Lemma 3, as y_A decreases from ∞ to 0, players' behaviour in equilibrium either I) remains at $T0$, II) changes from $T0$ to $T1_{(+)}$ or III) changes from $T0$ to $T1_{(+)}$ and then to $T2^{-}_{(+),(+-)}$. If $c < \widehat{c}$ it changes from $T0$ to $T2^{+}_{(++)}$ and then to $T1_{(+)}$ and then to $T2^{-}_{(+),(+-)}$. The result follows by Remark A1.

2.9.8.2 Moderate Leaning Towards Progressive DM

Suppose $c > \widehat{c}$, where $\widehat{c} = \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$. By Lemma 3, there exists some c such that as y_A decreases, players' behaviour changes from $T1_{(+)}$ to $T2^{-}_{(+),(-+)}$. The result follows by Remark A1.

Suppose $c \leq \widehat{c}$. By Lemma 3, the only change compared to observable testing is that $T2^{-}_{(+),(-+)}$ is played for y_A sufficiently low. By Remark A1, the DM's preferences are still single-peaked. Further, by

Lemma 3, the only change compared to observable testing occurs for $y_A < y_{DM}$. Hence, the argument in Section 2.9.7.2 for why a more reluctant advisor may be strictly preferred to an unbiased still applies.

2.9.8.3 Moderate Leaning Towards Conservative DM

Suppose $c > \widehat{c}$, where $\widehat{c} = \overline{c}^{+}_{12}\left(\frac{y_A}{x_A}\right) = \overline{c}^{R}_{02}\left(\frac{y_A}{x_A}\right)$. By Lemma 3, there exists some c such that as y_A decreases, players' behaviour changes from $T1_{(+)}$ to $T0$. The result follows by Remark A1.

Suppose $c \leq \widehat{c}$. By Lemma 3, as y_A decreases from ∞ to 0, players' behaviour changes from $T0$ to $T2^{+}_{(++)}$ and then to $T1_{(+)}$ and then to $T2^{+}_{(++)}$ at any given cost. Hence, by Remark 1A, preferences are no longer single peaked. Further, by Lemma 3, the only change compared to observable testing occurs for $y_A < y_{DM}$. Hence, the argument in Section 2.9.7.2 for why a more reluctant advisor may be strictly preferred to an unbiased advisor still applies.

2.9.8.4 Conservative DM

The result follows since, by Lemma 3, players' behaviour is independent of whether testing is observable or hidden.

Bibliography

A. Abdulkadiroglu and T. Sönmez. School choice: A Mechanism Design Approach. *The American Economic Review*, 93(3):729–747, 2003.

R. Allen, S. Burgess, and L. McKenna. School Performance and Parental Choice of School: Secondary Data Analysis. 2014.

L. Barseghyan, D. Clark, and S. Coate. Public School Choice: An Economic Analysis. Working Paper 20701, National Bureau of Economic Research, November 2014.

P. Bayer and R. McMillan. Choice and Competition in Local Education Markets. Working Paper 11802, National Bureau of Economic Research, November 2005.

C. Belfield and H. Levin. The Effects of Competition between Schools on Educational Outcomes: A Review for the United States. *Review of Educational Research*, 72(2):279–341, 2002.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

S. Black. Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics*, 114(2):577–599, 1999.

S. Burgess, A. Briggs, B. McConnell, and H. Slater. *School choice in England: Background facts*. Centre for Market and Public Organisation, University of Bristol Bristol, 2006.

S. Burgess, E. Greaves, A. Vignoles, and D. Wilson. What Parents Want: School Preferences and School Choice. *The Economic Journal*, 125(587):1262–1289, 2015.

S. Burgess, R. Metcalfe, and S. Sadoff. Using Behaviour Incentives to Improve Performance on High Stakes Tests: Evidence from a Field Experiment. *Unpublished manuscript*, 2016.

M. Burke and T. Sass. Classroom Peer Effects and Student Achievement. *Journal of Labor Economics*, 31(1):51–82, 2013.

S. Callander and J. Hörner. The Wisdom of the Minority. *Journal of Economic Theory*, 144(4):1421–1439, 2009.

C. Calsamiglia and M. Güell. The Illusion of School Choice: Empirical Evidence from Barcelona. 2014.

K. Carman and L. Zhang. Classroom Peer Effects and Academic Achievement: Evidence from a Chinese Middle School. *China Economic Review*, 23(2):223 – 237, 2012.

A. Chandra, A. Finkelstein, A. Sacarny, and C. Syverson. Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector. *American Economic Review*, 106(8):2110–44, 2016.

Y.-K. Che and N. Kartik. Opinions as Incentives. *Journal of Political Economy*, 117(5):815–860, 2009.

J. Coldron, E. Tanner, S. Finch, L. Shipton, C. Wolstenholme, B. Willis, S. Demack, and B. Stiell. Secondary School Admissions. *DCSF research report*, 2008.

D. Cutler, R. Ilckman, and M. Landrum. The Role of Information in Medical Markets: An Analysis of Publicly Reported Outcomes in Cardiac Surgery. *American Economic Review*, 94(2):342–346, 2004.

G. De Fraja and P. Landeras. Could do better: The Effectiveness of Incentives and Competition in Schools. *Journal of Public Economics*, 90(1):189–213, 2006.

G. De Giorgi, M. Pellizzari, and S. Redaelli. Identification of Social Interactions through Partially Overlapping Peer Groups. *American Economic Journal: Applied Economics*, 2(2):241–75, 2010.

D. Dranove, D. Kessler, M. McClellan, and M. Satterthwaite. Is More Information Better? The Effects of "Report Cards" on Health Care Providers. *Journal of Political Economy*, 111(3):555–588, 2003.

M. Felgenhauer and E. Schulte. Strategic Private Experimentation. *American Economic Journal: Microeconomics*, 6(4):74–105, 2014.

M. Fishman and K. Hagerty. The Optimal Amount of Discretion to allow in Disclosure. *The Quarterly Journal of Economics*, 105(2):427–444, 1990.

A. Gavazza and A. Lizzeri. The Perils of Transparency in Bureaucracies. *American Economic Review*, 97(2):300–305, 2007.

A. Gavazza and A. Lizzeri. Transparency and Economic Policy. *Review of Economic Studies*, 76(3): 1023–1048, 2009.

M. Gaynor, C. Propper, and S. Seiler. Free to choose? Reform and Demand Response in the English National Health Service. Technical report, National Bureau of Economic Research, 2012.

M. Gaynor, R. Moreno-Serra, and C. Propper. Death by Market Power: Reform, Competition, and Patient Outcomes in the National Health Service. *American Economic Journal: Economic Policy*, 5 (4):134–66, 2013.

Dino Gerardi and Leeat Yariv. Costly expertise. *The American Economic Review*, 98(2):187–193, 2008.

S. Gibbons, S. Machin, and O. Silva. Choice, Competition and Pupil Achievement. IZA Discussion Papers 2214, Institute for the Study of Labor (IZA), 2006.

J. Greene, T. Loveless, W B. MacLeod, T. Nechyba, P. Peterson, M. Rosenthal, and G. Whitehurst. Expanding Choice in Elementary and Secondary Education. *A Report on Rethinking the Federal Role in Education, Brown Center on Education Policy at Brookings*, 2, 2010.

J. Hastings and J. Weinstein. Information, School Choice, and Academic Achievement: Evidence from Two Experiments. NBER Working Papers 13623, National Bureau of Economic Research, Inc, November 2007.

E. Henry. Strategic Disclosure of Research Results: The Cost of Proving your Honesty. *The Economic Journal*, 119(539):1036–1064, 2009.

C. Hoxby. Peer Effects in the Classroom: Learning from Gender and Race Variation. Technical report, National Bureau of Economic Research, 2000.

S. Imberman and M. Lovenheim. Does the Market Value Value-Added? Evidence from Housing Prices After a Public Release of School and Teacher Value-Added. Working Paper 19157, National Bureau of Economic Research, June 2013.

T. Kane, D. Staiger, D. Grissmer, and H. Ladd. Volatility in School Test Scores: Implications for Test-Based Accountability Systems. *Brookings Papers on Education Policy*, (5):235–283, 2002.

D. Kessler and M. McClellan. Is Hospital Competition Socially Wasteful? *The Quarterly Journal of Economics*, 115(2):577–615, 2000.

V. Lavy and A. Schlosser. Mechanisms and Impacts of Gender Peer Effects at School. *American Economic Journal: Applied Economics*, 3(2):1–33, 2011.

V. Lavy, O. Silva, and F. Weinhardt. The Good, the Bad, and the Average: Evidence on Ability Peer Effects in Schools. *Journal of Labor Economics*, 30(2):367 – 414, 2012.

I. Lobel, D. Acemoglu, M. Dahleh, and A. Ozdaglar. Preliminary Results on Social Learning with Partial Observations. In *Proceedings of the 2nd International Conference on Performance Evaluation Methodologies and Tools*, 2007.

C. Manski. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*, 60(3):531–542, 1993.

M. A. Meyer. Learning from Coarse Information: Biased Contests and Career Profiles. *The Review of Economic Studies*, 58(1):15–41, 1991.

P. Milgrom. Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2):380–391, 1981.

P. Milgrom and J. Roberts. Comparing Equilibria. *The American Economic Review*, 84(3):441–459, 1994.

J. Rothstein. Good Principals or Good Peers? Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition among Jurisdictions. *American Economic Review*, 96(4):1333–1350, 2006.

D. Whitmore. Resource and Peer Impacts on Girls' Academic Achievement: Evidence from a Randomized Experiment. *American Economic Review*, 95(2):199–203, 2005.

D. Wilson and A. Piebalga. Accurate Performance Measure but Meaningless Ranking Exercise? An Analysis of the English School League Tables. Technical report, Department of Economics, University of Bristol, UK, 2008.