

Essays on Microeconomic Theory

Inauguraldissertation zur Erlangung des Doktorgrades
der Wirtschafts- und Sozialwissenschaftlichen Fakultät
der Universität zu Köln

2022

vorgelegt von

Marius Alexander Gramb, M.Sc.

aus Mönchengladbach

Referent: Prof. Dr. Christoph Schottmüller

Korreferent: Prof. Dr. Alexander Westkamp

Datum der Promotion: 17.05.2023

Acknowledgments

I am deeply indebted to my supervisor, Christoph Schottmüller, for his guidance and continuous support throughout the last years. His insightful comments on my work, his warm personality and almost daily conversations have been a constant source of inspiration and motivation. I was very fortunate to have Christoph as my supervisor. I would also like to thank my second supervisor, Alexander Westkamp, for introducing me to the fascinating world of matching theory. Chapter 3 of this thesis has benefited greatly from his advice and suggestions.

I am grateful to my co-author Julian Teichgräber for the fruitful collaboration and the good time during my research stay in Zurich. I thank Yiqiu Chen, Markus Möller and Lennart Struth for inspiring conversations about common research interests and for many useful comments and suggestions on my projects. I would also like to thank many more fellow doctoral students and colleagues for making my time as a researcher in Cologne such an enjoyable experience, in particular Viola Ackfeld, Emmanuele Bobbio, Simon Brandkamp, Kevin Breuer, Rebekka Cordes, Max R.P. Grossmann, Susanna Grundmann, Kiryl Khalmetski, Melisa Kurtis, Mark Marner-Hausen, Kirsten Marx, Yero Ndiaye, Max Thon and Tim Umbach.

Last but not least, I thank my family and friends for their constant emotional support throughout my studies, for always believing in me and, most importantly, for distracting me when needed.

Contents

1	Introduction	1
2	Game Preparation and Experience	4
2.1	Introduction	4
2.2	Related Literature	7
2.3	The Model	8
2.4	Analysis	11
2.5	Extensions	17
2.6	Comparison to Empirical Data: Opening Choices of Professional Chess Players	23
2.7	Conclusion	26
A	Proofs	28
B	Elements of Evolutionary Game Theory	46
C	Data Collection	50
3	Congestion and Market Thickness in Decentralized Matching Mar- kets	52
3.1	Introduction	52
3.2	Related Literature	55
3.3	The Model	56
3.4	An illustrative example	59
3.5	Myopic Strategy: Analysis	61
3.6	Strategic Targeting: Bayesian Nash Equilibrium Analysis	61
3.7	Impact of Market Thickness: Comparative Statics	67
3.8	Conclusion	72
A	Proofs	74

4	Anonymous or personal?	
	A simple model of repeated personalized advice	79
4.1	Introduction	79
4.2	Related literature	82
4.3	Model	84
4.4	Analysis	88
4.5	Welfare dynamics and anonymization	97
4.6	Experimental Design and Results	102
4.7	Discussion	105
4.8	Conclusion	107
A	Proofs	108
B	Experiment Instructions	120
C	Additional Results and Robustness Checks	128
	References	132

Chapter 1

Introduction

*I think game theory creates ideas
that are important in solving and
approaching conflict in general.*

Robert Aumann, Nobel Laureate

This dissertation consists of three essays in the field of Microeconomic Theory. In each chapter, I use a game-theoretic model to study strategic behavior in a complex economic situation. In each of these situations, there is a conflict of interest between the agents and their actions are interdependent. The theoretical findings in Chapters 2 and 4 are supported by empirical analyses based on observational data and a laboratory experiment. In Chapter 2, I examine the role of game preparation and experience in games commonly played by professional players, in which the opponent's prior behavior is observable. Chapter 3 deals with the strategic behavior of firms who make job offers to previously interviewed workers in the labor market. Chapter 4 analyzes a repeated adviser-consumer relationship in which the adviser can receive a bonus if the consumer takes certain actions.

The following paragraphs briefly summarize each chapter. Chapter 2 is titled "Game Preparation and Experience" and is single authored.¹ In this paper, I model game preparation and experience in games where players can observe their opponent's past plays. In these situations, players face a tradeoff between specialization and unpredictability. On the one hand, they can choose the same action every time they play the game. This gives them a high experience level in this action but has the disadvantage of being very predictable. On the other hand,

¹This chapter contains the current version of the working paper Gramb (2022).

players can choose to mix uniformly between all actions. This makes it harder for their opponent to prepare since these players are highly unpredictable. However, they also have less experience in each action they play. To model experience evolution, the stage game is infinitely repeated and one player's experience changes after each instance depending on the action taken in the stage game. This makes the overall game a dynamic game. The probability of winning a stage game after choosing a particular option increases with the player's experience in that option and decreases with the opponent's preparation time for that option. I find that both specialization and unpredictability are Nash equilibria of the stage game. To reduce this multiplicity of Nash equilibria, I apply methods from evolutionary game theory to the dynamic game. These methods are used to derive dynamic stability properties of the two Nash equilibrium behaviors specialization and unpredictability. The result is that specialization is always dynamically stable while unpredictability is sometimes not. This finding is also supported empirically by studying the opening choices of professional chess players which differ significantly from uniform mixing. The statification technique introduced to study dynamic games using evolutionary game theory methods can be applied to a large class of dynamic games. In this sense, the paper also contributes to the theory of dynamic games in general.

Chapter 3 is titled "Congestion and Market Thickness in Decentralized Matching Markets" and is joint work with Julian Teichgräber.² In this paper, we study congested decentralized matching markets and the impact of market thickness on market outcomes. The main application are labor markets, where many firms interview a similar set of workers and strategically make a job offer to one of these workers. We derive equilibrium strategies for each firm conditional on its rank among all firms and overall market thickness. We find that it is often not optimal for many firms to make an offer to the best observed worker, as the probability of acceptance may be low. In addition, we study how market outcomes change as market thickness increases. All firms and good workers lose when the market becomes thicker, while only low-skill workers benefit. In a similar vein, all firms and good workers would support a centralization of the market with an accompanying assortative matching, while only low-skill workers are in favor of the congested market structures in a decentralized labor market. Since market

²This chapter contains the current version of the working paper Gramb and Teichgräber (2023). Both authors developed the idea and the model together and derived most results together. Marius Gramb worked on the results about market thickness and finished the final draft of the paper.

thickness can be designed by different policy measures, our results have policy implications. By implementing a thicker or less thick market, policy makers can make different market participants better or worse off, depending on their policy goals. Moreover, our result on the support of a centralized market explains the formation of clearinghouses that reduce congestion problems in thick matching markets.

Chapter 4 is titled “Anonymous or personal? A simple model of repeated personalized advice” and is joint work with Christoph Schottmüller.³ In this paper, we study an adviser-consumer relationship in which the adviser receives a bonus if the consumer takes certain actions. In such situations, the adviser faces a strategic tradeoff: On the one hand, he wants to receive the highest possible bonus in every period. On the other hand, the adviser also wants to generate fitting advice, since the consumer might fire him if he gives too much bad advice. This would stop all future bonus payments. To study this strategic tradeoff, we introduce learning into the model. This means that the expert gets to know the consumer better whenever he gives advice that meets the consumer’s needs. This leads to a better signal quality in future periods and could incentivize the expert to give good advice even if it means not receiving a bonus. In the resulting dynamic game, we analyze two different types of equilibria. In *Markov equilibria*, the adviser exploits the consumer in the sense that the consumer does not benefit from the expert’s learning. In *m-equilibria*, the welfare generated by the learning opportunity is shared between consumer and expert. Complementary to our theoretical results, we conducted a laboratory experiment to see which equilibrium properties more closely resemble real-world behavior. We find that the learning opportunity provides an incentive for experts to give better advice. This explains why many people prefer personalized advice over anonymization. In the context of online advice, it justifies that many people use a personalized rather than an anonymized version of their preferred search engine.

³This chapter contains the current version of the working paper Gramb and Schottmüller (2022a). Christoph Schottmüller developed the idea and the model of the paper. Both authors derived several theoretical results. Marius Gramb programmed, conducted and evaluated the laboratory experiment. Both authors finished the final draft together.

Chapter 2

Game Preparation and Experience*

2.1 Introduction

In many situations, we anticipate playing a game against an opponent and we can also observe how she has previously decided in this kind of game. We are then faced with the problem of how to use our preparation time to prepare for various alternatives that are likely to occur. For example, consider a professor who is designing a final exam. A student will usually be able to see the most recent exams for that course and can infer the lecturer's preferences for what topics should be asked or emphasized. The lecturer, in turn, is well aware of these givens and might weigh the cost of designing entirely new assignments against the benefit of surprising the student in the exam. In other words, the lecturer faces a tradeoff between specialization and unpredictability: Specialization in one topic makes it easier to design new exercises based on a high level of experience but it comes at the cost of predictability. However, in order to be unpredictable, the professor has to use exercises from different topics which makes it more time-consuming to create the exam. Another example of this tradeoff is sports teams facing each other in a major competition and preparing for the opponent's preferred plays. A prominent instance of this is the Super Bowl, the annual finale of the American Football League. Each year, many resources are expended by the two opposing teams to analyze the opponents' offensive as well as defensive plays and determine

*I am grateful to Christoph Schottmüller, Yiqiu Chen, Markus Möller, Lennart Struth and numerous seminar and conference participants for valuable comments and suggestions. The remaining errors are my own.

which strategies offer the highest probability of victory. Some more economical applications of our framework include defending against terrorism¹ or choosing the agendas of political parties in an election campaign².

Finally, the same applies to chess opening preparation at the professional level. All the games of a professional chess player are publicly available and their opponents can see which openings they prefer to play and what their secondary weapons might be. This implies that also in this case, each professional must weigh the benefit of surprising the opponent by playing an opening different from his usual repertoire against the disutility from playing an opening in which he has significantly less experience.

In this paper, we model these situations as a dynamic two-player game where one player decides between different options while she is endowed with an experience vector reflecting the proportions with which she played these options in the past. The other player simultaneously decides how to allocate his preparation time among the alternatives, knowing his opponent's experience vector. We allow this experience vector to change over time based on the strategies chosen by the first player. The payoff of each player can be interpreted as their probability of winning the game, which naturally depends on both the first player's experience in the chosen option and the preparation time the second player used to prepare for that option.

In the framework described above, we are interested in answering several questions. The most obvious question is how to use one's preparation time most efficiently to prepare optimally for the opponent. Another one relates to a more meta-game theoretic standpoint: How broad should one's repertoire of plays be in order to make it difficult for the opponent to prepare and still retain a sufficiently high probability of winning? The underlying assumption in these situations is always that past plays also reflect the experience one has in playing those options and that more experience results in a higher probability of winning once the option in question is played. Note that a crucial assumption to make this framework relevant is that the opponents are of approximately equal strength. If one adversary is significantly better than the other, the role of preparation might become less pronounced.

¹Terrorists strike in places where security measures are low, but these measures are allocated based on an assessment of dangers.

²Different parties may have more experience or credibility in certain areas, but they must also tailor their agenda to the expected agendas of the opposing parties.

To study the above questions, we will step away from standard game theory and mainly adopt an evolutionary game theory (EGT) approach. While the Nash equilibrium is a useful tool to detect statically stable strategy combinations in a game, it does not predict which of these states is going to be reached in case of multiplicity. With EGT methods, however, it is possible to distinguish between different equilibria by evaluating their dynamic stability. In other words, when society is close to a statically stable state, these methods assess whether it will converge to this state or move away from it. In this way, we can address the issue of multiple Nash equilibria in our model.³ An intuitive justification for the use of an evolutionary approach might be that good preparation strategies as well as successful plays are passed from one coach to the next or from father to son while the worse strategies will become extinct quickly.⁴

The main finding of this paper is that focusing on specialization is more stable than focusing on unpredictability. Specifically, we show that being a predictable expert who always employs the same strategy is a stable outcome, while mixing uniformly between different alternatives is sometimes not. Intuitively, this is due to different incentives to change behavior in situations that are close to perfect specialization or uniform mixing. If a player's playing behavior has been close to uniform mixing for several rounds, then her opponent's optimal preparation behavior will also be close to uniform mixing. In this case, however, the player has an incentive to focus exclusively on one of the alternatives if this leads to a rapid increase in her experience in that strategy. In this way, she intends to take advantage of her opponents who are still preparing for both options equally. On the other hand, if the same player's playing behavior is close to perfect specialization, she has no incentive to play the less preferred option more often, since this leads to a lower probability of winning the game due to her limited experience in that option.

The rest of this paper is organized as follows: Section 2.2 discusses related literature. Section 2.3 presents the model of the paper and the analysis is performed in Section 2.4. Standard definitions and results from evolutionary game theory that are needed for the analysis are reviewed in Appendix B. Two natural extensions

³Equilibrium refinements such as perfect equilibrium or proper equilibrium are not able to eliminate this multiplicity of Nash equilibria.

⁴Concretely, a coach could pass on information about which kinds of plays worked well for the team and also if they were more successful with a certain preparation behavior against their opponents. Likewise, the father will pass on important life advice from his experience to his son and discourage certain approaches that did not work for him.

are studied in Section 2.5. Section 2.6 compares the results with the preparation behavior of professional chess players. Section 2.7 concludes.

2.2 Related Literature

The game studied in this paper resembles the well-known Colonel Blotto game introduced by Borel (1921). The main difference is that in our setting, there is only one “battlefield” which is chosen by player A , while player B hopes to have placed a large portion of his resources on the battlefield that is eventually chosen. Golman and Page (2009) generalize the Colonel Blotto game to a class of General Blotto games and find that under certain conditions, resources are allocated equally across all fronts in equilibrium. While this behavior is mirrored by player B in some stable states of our game, there is still a key difference in that only one battlefield is ultimately chosen and both players know player A ’s preferences over battlefields through her experience vector. In Hernández and Zanette (2013), the Colonel Blotto game with two types of players is analyzed within an evolutionary game theory framework and it is simulated which equilibria are reached depending on the starting conditions of the population. While we adopt a similar approach, the game we study changes continuously due to the possible change in player A ’s experience after each instance of the game.

Strategically, the game we study is also similar to the Matching Pennies game since player A attempts to mismatch player B ’s preparation while player B tries to match player A ’s action in his preparation. The dynamics of a repeated Matching Pennies game were studied by Becker et al. (2007) using different learning dynamics than we use in this paper. They find that there is no asymptotically stable state, but rather that play evolves in orbits around the unique equilibrium point. In our setting, the existence of asymptotically stable states is due to the experience vector of player A , which makes the game dynamic and guarantees that past and present plays can be consistent in certain situations such that we can expect convergence to one of the steady states of the system.

In this paper, we find that the dynamic game we study admits stable steady states that might not be expected in the underlying static game. The relationship between steady states of a dynamical system and the Nash equilibrium or other equilibrium concepts was studied also by Samuelson (1988), Mukhopadhyay and Chakraborty (2020), Juul et al. (2013) and Cheng and Yu (2018).

Gong et al. (2018) are probably closest to our paper methodologically since they study evolutionary dynamics of a dynamic game in general, for a particular form of *environmental feedback* that captures how the game changes depending on the population shares choosing different strategies. However, the class of dynamic games that can be captured by our approach is broader and Gong et al. also put more restrictions on the evolution rules of their game dynamics in order to obtain general results for games within this class. The technique we propose in Section 2.4.2 can be seen as a novel solution concept for dynamic games. Existing solution concepts for dynamic games are reviewed in Van Long (2010).

Since we mention the defense against terrorism as one possible application of our framework, this paper is also related to Cardoso and Diniz (2009), which studies this as well as protection against various other hazards from a game theoretical perspective. In Section 2.6, we study how professional chess players' opening choices are consistent with our game theoretical predictions. In the context of serves in tennis and penalty shoot-outs in soccer, a similar analysis is conducted in Walker and Wooders (2001), Anderson et al. (2021) and Palacios-Huerta (2003). Last but not least, evolutionary dynamics can also be interpreted as learning dynamics. An extensive treatment of learning in games can be found in Fudenberg and Levine (1998).

2.3 The Model

We consider a dynamic two-player game⁵ with infinitely many periods. In the stage game, player A (she) can choose between the two options 1 and 2 or play a mix between the two. Let s_i^t denote the probability that player A plays option i in stage t and write $s^t = (s_1^t, s_2^t)$. We assume each player's payoff in stage t depends on a state variable we call the *experience vector* $\alpha^t = (\alpha_1^t, \alpha_2^t)$ of player A , where α_i^t denotes the relative experience player A has in playing option i at stage t , with $\alpha_i^t \geq 0$ and $\alpha_1^t + \alpha_2^t = 1$. A player's payoff in the dynamic game is the sum of

⁵Note that this is different from a *dynamic game with (in)complete information* - a term sometimes used for extensive-form games. Also, (in)initely repeated games are dynamic games, but not vice versa. The difference is that stage game payoffs in a dynamic game can depend both on the players' actions and on a state variable that changes over time. A description of dynamic games can be found in Van Long (2010).

discounted stage game payoffs. The experience in future periods changes depending on today's play in the following way:

$$\alpha^{t+1} = (\kappa\alpha_1^t + (1 - \kappa)s_1^t, \kappa\alpha_2^t + (1 - \kappa)s_2^t), \quad (2.1)$$

where κ is an exogenously given parameter that reflects the relevance of experience to the specific stage game. We assume that κ is in $[0, 1)$, since $\kappa = 1$ would imply that there is no effect of recent plays on the experience vector and $\kappa = 0$ represents the extreme case in which the play in the last instance of the game completely determines the updated experience vector. This parameter κ is specific to the game in question. For instance, experience in certain playing options may be more important in the game of chess than in shooting a penalty in a soccer match.

Player B (he) can observe his opponent's experience vector and his strategy is given by a preparation vector $p^t = (p_1^t, p_2^t)$ with $p_1^t + p_2^t = 1$, indicating what fraction of time p_i^t is spent preparing for option i at stage t . The payoffs of the stage game - omitting the time index t - are assumed to depend linearly on experience and preparation time in the following way⁶:

$$u_A(s_1, p_1) = \sum_{i=1}^2 s_i \left(\frac{1}{2} + \frac{\alpha_i - p_i}{2} \right) = s_1 \left(\frac{1}{2} + \frac{\alpha_1 - p_1}{2} \right) + (1 - s_1) \frac{p_1 + \alpha_2}{2} \quad (2.2)$$

for player A and $u_B(s_1, p_1) = 1 - u_A(s_1, p_1)$ for player B . In this way, each stage game can naturally be interpreted as a zero-sum game (by subtracting $\frac{1}{2}$ from each entry in the payoff matrix). In the above version, the payoff corresponds to the probability of winning the stage game. If player A chooses a pure strategy ($s_i = 1$ for $i \in \{1, 2\}$), this payoff increases with experience α_i and decreases the more preparation time p_i player B spends on the chosen option. Note that the constant term of $\frac{1}{2}$ can be interpreted as the baseline probability of winning, reflecting the assumption that the players are of equal skill level. In Section 2.5.1, we will consider cases where the players' skill levels may differ to some extent. The normal form of the above game is shown in Table 2.1.

The following example illustrates the strategic tradeoff that player A faces.

⁶We denote payoffs depending on s_1 and p_1 only, since $s_2 = 1 - s_1$ and $p_2 = 1 - p_1$.

⁷We deviate from the usual notation of strategies denoted by s_1, s_2, \dots , since we are interested in the change in the probability that a particular option is played. Therefore, we denote pure strategies by the probabilities that the first option is played or prepared (e.g. $s_1 = 0$ or $s_1 = 1$) instead of s_1, s_2 and p_1, p_2 .

		Player B	
		$p_1 = 1$	$p_1 = 0$
Player A	$s_1 = 1$	$(\frac{\alpha_1}{2}, 1 - \frac{\alpha_1}{2})$	$(\frac{1+\alpha_1}{2}, \frac{1-\alpha_1}{2})$
	$s_1 = 0$	$(\frac{1+\alpha_2}{2}, \frac{1-\alpha_2}{2})$	$(\frac{\alpha_2}{2}, 1 - \frac{\alpha_2}{2})$

Table 2.1: Stage Game⁷

Example 1. Consider a setting in which $\alpha = (0.8, 0.2)$ and both players use pure strategies (surely playing one option and only preparing for one option, respectively). We note that $u_A(0, 0) = 0.1$, $u_A(0, 1) = 0.6$, $u_A(1, 1) = 0.4$ and $u_A(1, 0) = 0.9$. Hence, player A always prefers to meet player B unprepared over meeting him prepared (e.g., $u_A(0, 1) > u_A(1, 1)$), but fixing her opponent's level of preparation for the chosen option, she always prefers to play option 1 over option 2. More concretely, if for some reason player B can anticipate player A's chosen option and puts all of his preparation time into preparing for that option, then player A would always prefer to play option 1 because she has significantly more experience in that option. This captures the intuition that experience has some value for player A, but the effect of surprising the opponent is also valuable.

At this point, let us discuss how to interpret this model and how it relates to real-world interactions between professional players of a specific underlying game.

First, the two-player game in our model is not symmetric. In reality, both players would have some prior experience and both players would prepare simultaneously for the opponent. This means that, on the one hand, they decide which option to choose based on their own experience. On the other hand, they also allocate time shares to prepare for certain options based on their opponent's experience. The asymmetric game in our model allows us to disentangle these two strategic considerations.

Second, the model should be interpreted as a reduced form of a very complex game played in reality. This is also the reason why we use the term *option* instead of *action*: One option can be interpreted as a particular way of playing the game, which could subsume several different but similar actions or sequences of actions in the real-world game. For example, in the design of an exam task, an option might correspond to the choice of a particular topic and the student's success in this task will depend (in expectation) on how much preparation time was put into studying that specific topic. Similarly, a chess player might choose a particular first move in a chess game, after which the game quickly ramifies. However, both her experience

with and her opponent's preparation time for positions related to that specific first move are relevant in determining her probability of winning the game.

We conclude this section by discussing why it seems more appropriate to model relative experience than absolute experience levels. First, the strategic considerations of both players should be driven by the comparative experience advantage player A has in one option over the other. The absolute number of games played with each alternative is less relevant for these considerations. The second reason is that equal skill among players cannot be modeled using absolute experience but can be modeled with relative experience. The technical problem with absolute experience levels is that experience can tend to infinity for an infinite number of periods while player B 's preparation time is always a number in the compact set $[0, 1]$. Note, however, that uneven skill between players can also be modeled with relative experience, as we will show in Section 2.5.1.

2.4 Analysis

In this section, we will analyze the game presented in Section 2.3. To do so, we first compute the Nash equilibria of one independent stage game for a given experience vector α in Section 2.4.1. In Section 2.4.2, we introduce a novel approach that allows us to study the dynamic game (where α adapts in the course of the game) with standard evolutionary game theory tools. These tools as well as several central definitions are reviewed in Appendix B.

2.4.1 The static one-shot game

Before considering player A 's experience evolution in the dynamic game, it is useful to examine the static one-shot game, where the experience vector is exogenously given.

Proposition 1. *For a given experience vector $\alpha = (\alpha_1, \alpha_2)$ with $\alpha_1 \in (0, 1)$, the unique Nash equilibrium in mixed strategies of the stage game is given by $s = (\frac{1}{2}, \frac{1}{2})$ and $p = \alpha$. For $\alpha_1 \in \{0, 1\}$, $s = (\frac{1}{2}, \frac{1}{2})$ and $p = \alpha$ is still a Nash equilibrium and there is an additional Nash equilibrium in pure strategies given by $s = p = \alpha$.*

Proof. See Appendix A. □

This result gives two insights. First, when player A has experience in both options, we see that uniform mixing between these two options is her unique Nash

equilibrium strategy. In this Nash equilibrium, player B will prepare according to the experience vector by player A ($p = \alpha$), preparing more for the option in which A has more experience. Second, when player A is already perfectly specialized in one option ($\alpha_1 = 0$ or $\alpha_1 = 1$), both $s = (\frac{1}{2}, \frac{1}{2})$ and $s = \alpha$ are Nash equilibrium strategies of player A . Hence, she might go for specialization or unpredictability in this scenario. Let us assume for the moment that player A has experience in both options already and consider the dynamic game where α is allowed to adjust based on past play. In this game, unique Nash equilibrium behavior by both players at every stage would lead to a balanced experience structure ($\alpha = (\frac{1}{2}, \frac{1}{2})$) in the long run, due to A 's uniform mixing. This seems to suggest that all players in society should favor unpredictability over specialization unless they have been perfectly specialized to start with (and even then, they might go for uniform mixing instead of continuing to specialize in their favorite option). In the next subsection as well as in Section 2.5, we will see that the opposite is true: Among the two equilibrium behaviors specialization and unpredictability, specialization can be seen as more evolutionarily stable than unpredictability.

2.4.2 An evolutionary approach

In this section, we study whether either of the two Nash equilibria identified in Section 2.4.1 – specialization and unpredictability – is a more reasonable prediction of aggregate behavior in a society. To do so, let us first recall a microfoundation of Nash equilibrium given by Nash himself in his dissertation:

We shall now take up the “mass-action” interpretation of equilibrium points. [...] It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.

To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the “average playing” of the game involves n participants selected at random from the n populations, and that there is a stable average frequency with which each pure strategy is employed by the “average member” of the appropriate population. Nash (1950)

This "mass-action" interpretation can be modeled with population games. The underlying idea is that a continuum of individuals plays a particular game infinitely often and individuals revise their strategies based on their success in the game. Hence, it is explicitly modeled how and why a population reaches a particular equilibrium point. In this sense, one Nash equilibrium may have better properties in terms of dynamic stability than another. This is precisely what we will investigate in this section. To this end, we use evolutionary game theory and assess the dynamic stability of specialization and unpredictability using replicator dynamics. It should be noted that already von Neumann and Morgenstern (1944, 44) stated "most emphatically that [their] theory is thoroughly static. A dynamic theory would unquestionably be more complete and therefore preferable." The evolutionary framework also seems well grounded practically in the examples given in Section 2.1: These professionals typically play the same game very often over the course of their careers and it is natural to assume that over time, through trial and error, they get a good feel for which strategies work and which do not.

A potential limitation of the evolutionary approach is that players are not forward-looking. Concretely, this means that they do not consider the impact of their decisions on future experience vectors. However, it is questionable whether forward-looking behavior is strategically desirable in our model. This is because it is not clear why player A should prefer a specific experience vector over another in a future stage game. The opposing player B can always keep player A indifferent between her two options by choosing preparation shares equal to A 's experience vector α . Moreover, forward-looking behavior implies that player A values some later stage games more than earlier ones and modeling this fact would make the model less tractable. All in all, the evolutionary approach appears suitable to study the dynamic stability of the two Nash equilibria specialization and unpredictability.

Note that the two-player game we study is dynamic as the stage game changes in every instance (due to the adaptation of player A 's experience). The standard tools of evolutionary game theory assume a static stage game matrix and thus cannot capture our experience dynamics. This is why we introduce a novel approach, which we call *statification*, that allows us to transform our dynamic game into a repeated game with a static stage game and analyze it using the standard tools of evolutionary game theory. Specifically, we will use the eigenvalue technique to assess the (dynamic) stability of certain population states. The idea of our transformation is simple: We introduce a virtual third player whose strategies reflect the experience

vector of player A and who thus determines the dynamics of the game with his strategy choice. In this way, the dynamic two-player game is transformed to an infinitely repeated three-player game. More concretely, virtual player C has two options 1 and 2 available to him and his choice determines the payoff matrix for players A and B :

- If player C chooses option 1, the payoff-matrix is

		Player B	
		$p_1 = 1$	$p_1 = 0$
Player A	$s_1 = 1$	$(\frac{1}{2}, \frac{1}{2})$	$(1, 0)$
	$s_1 = 0$	$(\frac{1}{2}, \frac{1}{2})$	$(0, 1)$

- If player C chooses option 2, this yields the payoff-matrix

		Player B	
		$p_1 = 1$	$p_1 = 0$
Player A	$s_1 = 1$	$(0, 1)$	$(\frac{1}{2}, \frac{1}{2})$
	$s_1 = 0$	$(1, 0)$	$(\frac{1}{2}, \frac{1}{2})$

If we now assume that player C chooses option 1 with probability α_1 and option 2 with probability $\alpha_2 = 1 - \alpha_1$, then the expected payoffs of players A and B conditional on their chosen strategies are given by the matrix in Table 2.1. Player C , on the other hand, does not receive a payoff himself, but rather adjusts his behavior according to the strategy chosen by player A . Concretely, we assume a continuous version of (2.1) in the sense that the time derivative of α_1 is given by $\dot{\alpha}_1 = (s_1 - \alpha_1)(1 - \kappa)$. This describes the evolution of virtual player C 's strategy. Note that these deterministic dynamics still imply that player A influences her future experience vectors through her behavior. However, this influence is now endogenized through virtual player C and therefore the stage game that is played by players A , B and C is the same in every period. Using this, we can study a dynamical system in continuous time t , which is more convenient than discrete time as the stability properties of certain states can be derived from the differential equations defining this system. We denote a state in this system at time t by

$x^t = (\alpha_1^t, s_1^t, p_1^t) \in [0, 1]^3$. The dynamical system is $\dot{x} = V(x)$, where $V: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is given by⁸

$$V_1(x) = \dot{\alpha}_1 = (s_1 - \alpha_1)(1 - \kappa) \quad (2.3)$$

$$V_2(x) = \dot{s}_1 = s_1(1 - s_1)(\alpha_1 - p_1) \quad (2.4)$$

$$V_3(x) = \dot{p}_1 = p_1(1 - p_1)(s_1 - \frac{1}{2}). \quad (2.5)$$

In the above system, equations (2.4) and (2.5) represent the replicator dynamics of s_1 and p_1 , the calculation of which can be found in Appendix A. Equation (2.3) describes the dynamics of the experience vector in our model. However, since these are different from the replicator dynamics and the associated growth-rate function is not Lipschitz continuous, we cannot apply Theorem 4 or Proposition 2 from Appendix B. Nevertheless, it is useful to keep them in mind as benchmark results. The following results identify the steady states of the system and assess their stability properties.

Lemma 1. *The steady states of the dynamical system are $(0, 0, 0)$, $(0, 0, 1)$, $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, $(1, 1, 0)$ and $(1, 1, 1)$.*

Proof. The steady states of the dynamical system are exactly the values of x for which $V(x) = 0$ holds. Since $\kappa < 1$, we see that s_1 and α_1 will always be equal in a steady state and solving $V_2(x) = 0$ and $V_3(x) = 0$ yields the five points mentioned. \square

Theorem 1. *The points $(1, 1, 1)$ and $(0, 0, 0)$ are asymptotically stable steady states of the system. The point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is an asymptotically stable steady state for $\kappa > \frac{3}{4}$ and a saddle point for $\kappa < \frac{3}{4}$. The other steady states $(0, 0, 1)$ and $(1, 1, 0)$ are saddle points.*

Proof. See Appendix A. \square

Interestingly, we find that the point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is not always asymptotically stable. In particular, it is only a saddle point when the role of experience κ is too small. The reason for this is the following: When κ is small, it is easier to build up a good experience level in an option quickly by playing it once or a few consecutive times. This way, player A has no interest in playing both options with equal share

⁸For simplicity, we omit the time index t from now on. After all, we are not interested in specific time periods, but in the change in overall strategies and experience over time.

again after playing one option several times in a row. For instance, suppose that $\kappa = 0.6$ and the population is in state $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Now, if player A plays option 1 in three consecutive (discrete) instances of the game, this already results in an experience vector of $\alpha = (0.892, 0.108)$. In this way, player A has a high incentive to continue playing option 1, since her experience in this option is significantly higher than her experience in the alternative option 2. However, if κ is higher, player A is more bound by her experience structure, since focusing on one option for a few consecutive rounds does not change the experience vector very much. For example, suppose that $\kappa = 0.95$ and the population is back in state $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. In this case, playing option 1 three times in a row leads to an experience vector of $\alpha \approx (0.571, 0.429)$ for player A . The experience difference between the two options is much smaller than before. This explains intuitively why players do not have much incentive to move away from the point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ when κ is high. Note that being a predictable specialist (only playing one option in every instance of the game) is always asymptotically stable, regardless of the value of κ . This might seem counter-intuitive at first glance, since in practice people are sometimes criticized for using the same game strategy over and over again. Moreover, we saw in Proposition 1 that $s = (\frac{1}{2}, \frac{1}{2})$ is the unique Nash equilibrium strategy for player A even for values of α_1 close to 1. One way to interpret the finding in Theorem 1 is that evolution will favor specialization and expertise over unpredictability in these kinds of repeated games where experience and the opponent's preparation time influence the player's probability of success.

2.4.3 The statification technique

We end this section with some remarks on the statification technique and what types of dynamic games can be transformed with it. To simplify the exposition, we assume a population game with two populations, where each population can choose from two different actions. Extending this to more populations or more actions is straightforward. Such a game is generally given by two payoff matrices $U_A = (a_{ij})$ and $U_B = (b_{ij})$, which capture the payoffs of population 1 and population 2, respectively. From these two matrices, we can construct a matrix $C = (c_{ij})$ whose entries $c_{ij} = (a_{ij}, b_{ij})$ are tuples consisting of the payoffs of the two populations.

Now, if we let a virtual player choose between two such matrices C_1 and C_2 , the resulting set of possible payoff matrices is given by

$$\{\alpha C_1 + (1 - \alpha)C_2 = C_2 + \alpha(C_1 - C_2) | \alpha \in [0, 1]\}.$$

This class of two-population games is exactly the class that Gong et al. (2018) can capture with their *environmental feedback* framework. This means that the class of dynamic games that can be transformed into static games using the statification technique is broader, since we could allow the virtual player to choose from multiple options or construct more evolved types of dynamics by introducing multiple virtual players (in this case, the payoff matrices would depend on the combination of virtual players' choices and all polynomial dependencies would be possible). Moreover, the evolution rules for the parameters determining the dynamics can be given by any differentiable function and need not be replicator dynamics (which is the assumption in Gong et al. (2018)).

2.5 Extensions

In this section, we consider two natural ways of extending our model from Section 2.3. First, in Section 2.5.1, we relax the assumption that both players have the same playing strength. Then, in Section 2.5.2, we consider a model in which the number of options both players can choose or prepare for is increased from two to three.

2.5.1 Uneven Skill Distributions

A natural way to extend the analysis of Section 2.4 is to admit different *baseline probabilities of winning* reflecting potential skill differences between the two players. The baseline probability of winning denotes a player's winning probability when player B 's preparation vector equals player A 's experience vector. The assumption in Section 2.3 was that it is $\frac{1}{2}$ for both players. In order to analyse uneven skill distributions, we denote the players' skill levels by q_A and q_B . We define the baseline probability of player i winning as $\frac{q_i}{q_A + q_B}$. This probability is then converted to the actual probability of winning the game by taking into account the experience as well as the choice of player A and the preparation vector of player B in exactly the same way as in (2.2). The only thing to note here is that the resulting numbers

may be less than 0 or greater than 1 in case the skill difference is substantial and the stronger player anticipated the weaker player's action. In this case, we will correct the numbers to 0 or 1, respectively, in order to be able to interpret the payoff as the winning probability of the game.⁹

Before proceeding, a comment on the difference between a player's skill level and her experience vector is in order. On the one hand, the skill level is meant to measure the player's general ability to play the game in question. Her experience vector, on the other hand, provides information about how often she has played various options in the past. Imagine a highly skilled player has an experience level of only 0.1 in some option and chooses that option against a less skilled player. If the skill difference is large enough, this can still result in the more skilled player defeating the less skilled player, especially if this option has not been prepared much by the weaker player. Thus, the experience vector measures the extent to which a player with a given skill level plays one option relatively better than another option. The skill level, however, measures overall playing strength compared to other players.

Without loss of generality, we will normalize the skill level of the stronger player to 1 such that the skill level of the weaker player is a number $q_i \in [0, 1)$. Specifically, we can write down the payoff matrices of the stage game for a given experience vector α in both cases where either player A or player B is the stronger player:

- i) Player A is the stronger player. In this case, let $q_A = 1$ and $q_B \in [0, 1)$ be the skill level of player B . For a given experience vector α , the game matrix is given in Table 2.2.

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(\frac{1}{1+q_B} - \frac{1}{2} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{1}{2} - \frac{\alpha_1}{2})$	$(\min\{1, \frac{1}{1+q_B} + \frac{\alpha_1}{2}\}, \max\{0, \frac{q_B}{1+q_B} - \frac{\alpha_1}{2}\})$
$s_1 = 0$	$(\min\{1, \frac{1}{1+q_B} + \frac{1-\alpha_1}{2}\}, \max\{0, \frac{q_B}{1+q_B} - \frac{1-\alpha_1}{2}\})$	$(\frac{1}{1+q_B} - \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{\alpha_1}{2})$

Table 2.2: Player A is the stronger player

- ii) Player B is the stronger player. In this case, let $q_B = 1$ and $q_A \in [0, 1)$ be the skill level of player A . For a given experience vector α , the game matrix is given in Table 2.3.

⁹Of course, winning probabilities of exactly 0 or 1 are not realistic. On the other hand, they should be reasonably close to these numbers when a more skilled player meets a significantly less skilled player who is unprepared for the chosen option.

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(\max\{0, \frac{q_A}{1+q_A} + \frac{\alpha_1-1}{2}\}, \min\{1, \frac{1}{1+q_A} + \frac{1-\alpha_1}{2}\})$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1}{2}, \frac{1}{1+q_A} - \frac{\alpha_1}{2})$
$s_1 = 0$	$(\frac{q_A}{1+q_A} + \frac{1-\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1-1}{2})$	$(\max\{0, \frac{q_A}{1+q_A} - \frac{\alpha_1}{2}\}, \min\{1, \frac{1}{1+q_A} + \frac{\alpha_1}{2}\})$

Table 2.3: Player B is the stronger player

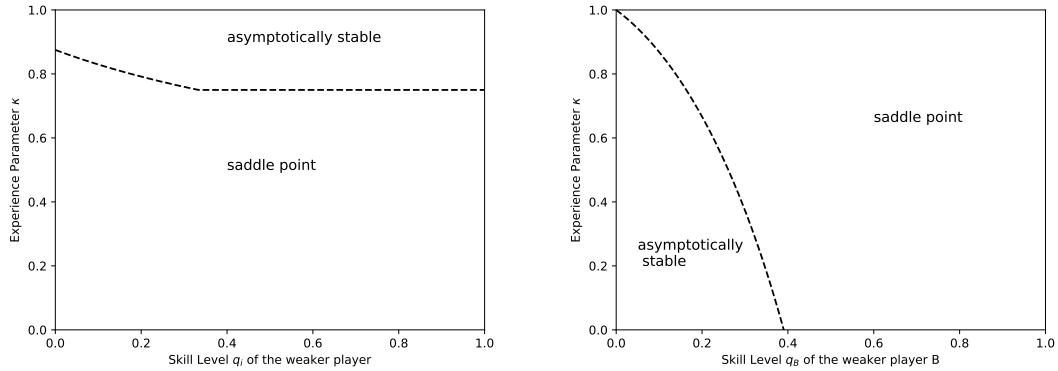
After determining the values of the maxima and minima as a function of certain parameter regions, it is possible to extend the stability analysis from Section 2.4 to this case of uneven skill distributions. It turns out that the results from Section 2.4 are largely robust to this change:

Theorem 2. *In the case of uneven skill distributions ($q_A \neq q_B$), let us assume that the skill level of the stronger player is normalized to 1 and the skill level of the weaker player i is given by $q_i \in [0, 1)$. The stability properties of the steady states of the dynamical system are as follows:*

- The points $(0, 0, 0)$ and $(1, 1, 1)$ are always asymptotically stable.
- The points $(0, 0, 1)$ and $(1, 1, 0)$ are always saddle points when player B is the stronger player. When player A is the stronger player, they are saddle points for $\kappa > \frac{1-q_B-4q_B^2}{1-q_B^2}$ and they are asymptotically stable for $\kappa < \frac{1-q_B-4q_B^2}{1-q_B^2}$.
- The stability of the point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ depends on the skill difference between the players: If $q_i > \frac{1}{3}$, it is asymptotically stable for $\kappa > \frac{3}{4}$ and a saddle point for $\kappa < \frac{3}{4}$. If $q_i < \frac{1}{3}$, it is asymptotically stable for $\kappa > \frac{21q_i^2+58q_i+21}{8(7q_i+3)(q_i+1)}$ and a saddle point for $\kappa < \frac{21q_i^2+58q_i+21}{8(7q_i+3)(q_i+1)}$.

Proof. See Appendix A. □

Comparing these results with Theorem 1, two interesting findings emerge. First, we see that the condition for $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ to be stable is strengthened when the skill difference between players increases. This seems intuitive: If player A is a lot stronger than player B , she might not care much about building up a high experience level in the option she plays, as long as playing that option (more often) comes as a surprise to player B . By surprising player B , we simply mean that player A tries to mismatch the preparation vector of player B . She could do so by playing an option that she has less experience in as player B might be less prepared for that option. In this case, she can rely on the large skill difference to outplay her opponent. This justifies that the point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ may not be stable even for higher values of κ if the skill difference is sufficiently large.



(a) For the stability of mixing,
independent of who is stronger

(b) For the stability of $(0,0,1)$ and $(1,1,0)$
when player A is stronger

Figure 2.1: Plots of the critical threshold values for κ

The second result, which seems surprising at first glance, is that the points $(0, 0, 1)$ and $(1, 1, 0)$ can actually be stable when player A is the stronger player. To understand this finding, suppose that the population is close to the state $(0, 0, 1)$ at $(\epsilon, \epsilon, 1 - \tilde{\epsilon})$, with $\epsilon, \tilde{\epsilon}$ positive but close to zero. This means that player A plays option 2 almost exclusively, which is also captured by her experience vector. Player B , on the other hand, prepares almost exclusively for option 1. The fact that the point $(0, 0, 1)$ is asymptotically stable means that for some small values of ϵ and $\tilde{\epsilon}$, the state will move even closer to $(0, 0, 1)$ starting from $(\epsilon, \epsilon, 1 - \tilde{\epsilon})$. The intuition for this is the following: Player B finds that he is significantly weaker than player A , perhaps even to the extent that he would certainly lose the game even if he prepared a bit more for option 2 and that option was eventually played. However, he also observes that player A might play option 1 in a tiny fraction of the cases and that she has little experience with this option. This means that preparing almost exclusively for this option can make up for the large skill difference and give player B a relatively high probability of winning the game in this particular case. The prospect of this event occurring motivates player B to increase his preparation share for option 1. Similarly, player A realizes that she is almost surely winning if she chooses option 2, which motivates her to play that option even more. This causes the population to move toward the state $(0, 0, 1)$ when κ is low relative to the skill difference between the players. The plots of the exact critical values for κ are shown in Figure 2.1.

2.5.2 Three different options

In this section, we study the effect of adding a third option to the choice set of players A and B . As a consequence, the experience vector of player A takes the form $\alpha = (\alpha_1, \alpha_2, 1 - \alpha_1 - \alpha_2)$, with $\alpha_i \geq 0$ for $i = 1, 2$ and $\alpha_1 + \alpha_2 \leq 1$. A population state in this system is given by a vector $x = (\alpha_1, \alpha_2, s_1, s_2, p_1, p_2) \in \mathbb{R}^6$, since the shares of choosing option 3 or preparing for option 3 can be given by $s_3 = 1 - s_1 - s_2$ and $p_3 = 1 - p_1 - p_2$, respectively. The game matrix for a stage game with given experience vector α is shown in Table 2.4.

		Player B		
		$p_1 = 1$	$p_2 = 1$	$1 - p_1 - p_2 = 1$
Player A	$s_1 = 1$	$(\frac{\alpha_1}{2}, 1 - \frac{\alpha_1}{2})$	$(\frac{1+\alpha_1}{2}, \frac{1-\alpha_1}{2})$	$(\frac{1+\alpha_1}{2}, \frac{1-\alpha_1}{2})$
	$s_2 = 1$	$(\frac{1+\alpha_2}{2}, \frac{1-\alpha_2}{2})$	$(\frac{\alpha_2}{2}, 1 - \frac{\alpha_2}{2})$	$(\frac{1+\alpha_2}{2}, \frac{1-\alpha_2}{2})$
	$1 - s_1 - s_2 = 1$	$(\frac{2-\alpha_1-\alpha_2}{2}, \frac{\alpha_1+\alpha_2}{2})$	$(\frac{2-\alpha_1-\alpha_2}{2}, \frac{\alpha_1+\alpha_2}{2})$	$(\frac{1-\alpha_1-\alpha_2}{2}, \frac{1+\alpha_1+\alpha_2}{2})$

Table 2.4: Stage Game with 3 options

In the following, we will use this game matrix and the associated selection dynamics to derive the steady states of the dynamical system with three options as well as their stability properties.

Lemma 2. *The steady states of the dynamical system with three options can be classified into five different categories i) through v). Within each category, an exemplary steady state or set of steady states is given, and the remaining steady states within that category differ only by relabeling the different options.*

i) *Uniform mixing between all options:* $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

ii) *Uniform mixing between two options with completely inconsistent¹⁰ preparation, e.g.* $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0)$

iii) *Uniform mixing between two options with consistent preparation, e.g.* $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$

iv) *Perfect specialization with completely inconsistent preparation, e.g.* $(1, 0, 1, 0, 0, p_2)$,
 $p_2 \in [0, 1]$

v) *Perfect specialization with consistent preparation, e.g.* $(1, 0, 1, 0, 1, 0)$

Proof. See Appendix A. □

¹⁰By *completely inconsistent* preparation, we mean that player B prepares exclusively for options in which player A has zero experience. *Consistent* preparation refers to the case $p = \alpha$.

Theorem 3. *Among the steady states presented in Lemma 2, only the ones associated to perfect specialization with consistent preparation (Category v)) are asymptotically stable steady states. The steady states in categories i) through iv) are all saddle points.*

Proof. See Appendix A. □

This result might seem surprising, especially when compared to Theorem 1. Even though uniform mixing between two options with consistent preparation by player B was an asymptotically stable steady state for high values of κ when there were only two alternatives, this state is reduced to a saddle point once the possibility of choosing a third alternative occurs. Uniform mixing between all three different alternatives (with consistent preparation) is not stable for any value of κ , either. The reason for this is the following: In a state close to uniform mixing between all three options, Player A has an incentive to play the least expected option a few times in a row to exploit the newly created gap between built-up experience in that option and the preparation time player B has spent on that option (based on his low expectations). Since player B will quickly adjust his expectations, player A could subsequently switch to another option using the same logic as before. This happens either continuously or until the experience in one option is so high that the state converges to perfect specialization (which is the only asymptotically stable state here). Simulations show that the system converges to one of the corner states corresponding to perfect specialization, even when the initial state is close to uniform mixing. An example of this can be seen in Figure 2.2.

The intuition for the instability of uniform mixing between two out of three options is similar: The presence of the third option, which player B does not expect, allows player A to increase her use of that option to her advantage for a few successive instances of the game. This leads to her playing profile approaching a mixture of all three options, or even a specialization in the third option.

A final point to notice is that perfect specialization in one option (with consistent preparation) continues to be an asymptotically stable state. This seems intuitive: Even with two options, none of the players wanted to shift attention from the frequently played or prepared option to the rarer one. Adding a third, rarely used option does not change this decision, since the incentives to choose this new option are the same as for the rare option that was already available before.




Figure 2.2: Simulation with $\kappa = 0.95$ and initial values $s_1 = p_1 = \alpha_1 = 0.35$, $s_2 = p_2 = \alpha_2 = 0.31$

2.6 Comparison to Empirical Data: Opening Choices of Professional Chess Players

In this section, we provide empirical evidence to support our theoretical results. To this end, we consider the opening choices of world-class professional chess players. The game of chess is well suited for these kinds of considerations as players' preparations can be explicitly observed in the game. Before a game, players usually prepare for a long time by checking their opponent's games and considering their own opening moves. If a chess player plays a certain move in the opening, you can be reasonably sure that this was already intended before the game (this typically holds for at least the first 5 to 10 or even 15 moves). Since all the moves are written down and games are stored in large databases, it is easy to study the top players' opening repertoires and determine how often they choose which opening variation.

In the following, we focus on the eight players who participated in the 2020 Candidates Tournament (which ended in 2021 due to the pandemic). The Candi-




Figure 2.3: Lorenz curve of top chess players' opening repertoires with Black after 1. e2-e4

dates Tournament is the second most important individual tournament in chess because its winner becomes the challenger to the reigning World Champion. For this reason, the players who qualify for this tournament are usually among the best in the world (excluding the World Champion himself). We look at the games played by these eight players in the two years leading up to this tournament (i.e., 2018 and 2019). In those games, we examine their replies with the black pieces to the most common move 1. e2-e4¹¹. The reason for doing this is the following: After 1. e2-e4, the opening system is in most cases pinned down by Black's reply on the first move and transitions to other openings are unlikely to occur later in the game. This is not the case with the other three main opening moves 1. d2-d4, 1. c2-c4 and 1. Ng1-f3, and there are many possible transitions where different initial moves eventually lead to the same position later on. This makes it difficult to analyze whether a player has played a different opening or only chosen a different path to reach the same opening and use his knowledge as well as his home preparation

¹¹This means that White starts by moving the pawn in front of the king two squares forward.

for this opening. The games included in the analysis were all played with classical time control (one game usually lasts between four and six hours and players play only one such game per day). Games with shorter time controls were excluded from the data set because the preparation times for these formats are much shorter and sometimes even non-existent (as many games against different opponents are played in one day with only short breaks in between). The total number of games remaining for each of these players over this two-year period ranges from 19 to 47 games. The exact procedure of data collection is described in Appendix C.

To get a first idea of whether players prefer to be an expert in one opening line by always sticking to it, or whether they prefer to mix between different openings to keep their opponents guessing, we use the Lorenz curve measuring the concentration of openings in their repertoire. If a player uses n different openings against 1.e2-e4, the curve plots the points $(i/n, s(i))$, for $i \in \{0, \dots, n\}$, where $s(i)$ is the share of games in the least used i openings of that player. Thus, if a player were to mix uniformly between several openings, the Lorenz curve would be equal to the 45-degree line. The more concentrated his repertoire is, the closer the curve is to the x -axis. The Lorenz curves for our sample of eight players and their games is depicted in Figure 2.3. Only one of these eight players played one reply exclusively. For this player, the point $(99/100, 0)$ was added to the Lorenz curve because otherwise it would give the impression that he mixed uniformly.¹²

Player	χ^2 statistic	df
Kirill Alekseenko	(15.31) ^{***}	2
Fabiano Caruana	(30.01) ^{***}	2
Anish Giri	(14.73) ^{**}	3
Alexander Grischuk	(8.91) ^{**}	1
Wang Hao	(22.49) ^{***}	3
Ding Liren	(28.49) ^{***}	1
Ian Nepomniachtchi	(24.14) ^{***}	3
Maxime Vachier-Lagrave	-	-

^{**} significant at the 1% level

^{***} significant at the 0.1% level

Table 2.5: Test for discrete uniform distribution of the black repertoires

¹²Of course, mixing uniformly and playing the same option exclusively are the same thing when there is only one option to choose from. However, we want to emphasize that this repertoire choice is highly concentrated, which is not visible in the graph without the added point.

We see that no one among these 8 players mixes uniformly. They tend to have one main weapon that they like to use, and beyond that they may have one or two secondary weapons that they use from time to time to surprise their opponents. To test whether this finding is significant, we perform a χ^2 test with the null hypothesis that each player mixes uniformly between all the openings he played in his games and that the openings chosen in each game are statistically independent. That is, we assume that the opening choices follow a discrete uniform distribution on all openings chosen with positive share in the games.¹³ The findings are summarized in Table 2.5. We see that all the players' opening repertoires differ significantly from uniform mixing while one player uses the same opening in every single game (Maxime Vachier-Lagrave, therefore no χ^2 test was performed for him). Although further tests could be conducted to study the difference in preparation against equally strong or significantly weaker opponents, these results seem to indicate that the goal of strong players is to be an expert in mainly one style of play. In other words, professional chess players tend to favor specialization over unpredictability. This is in line with the observation from Section 2.4 that states $(0, 0, 0)$ and $(1, 1, 1)$ are always asymptotically stable.

2.7 Conclusion

In this paper, we used tools known from evolutionary game theory to study the question of how to prepare for a game and how diverse one's repertoire of plays in such a game should be. We find that having a narrow repertoire and taking the same action exclusively whenever you play the game is always stable, even though this goes along with being predictable. Uniform mixing is player A 's equilibrium behavior in the unique interior Nash equilibrium in an associated static game. However, this behavior is not always dynamically stable. Looking at empirical data of professional chess players' opening choices seems to suggest that they do not believe in uniform mixing being an optimal strategy, either.

The inclusion of a virtual player to model a dynamic two-player game as a repeated three-player game with static stage game was immensely helpful as concepts known from evolutionary game theory could easily be applied. This might be a fruitful new technique to analyze dynamic games in economics and other

¹³Note that we could also assume that each player mixes between all the options chosen by the whole set of players. In this case, the results would be even more significant.

fields in cases where adaptive learning mechanisms are used to determine stable outcomes in a society.

Appendix

A Proofs

Proof of Proposition 1. The following table shows the payoffs of the stage game when both players only choose pure strategies.

		Player <i>B</i>	
		$p_1 = 1$	$p_1 = 0$
Player <i>A</i>	$s_1 = 1$	$\left(\frac{\alpha_1}{2}, 1 - \frac{\alpha_1}{2}\right)$	$\left(\frac{1+\alpha_1}{2}, \frac{1-\alpha_1}{2}\right)$
	$s_1 = 0$	$\left(\frac{1+\alpha_2}{2}, \frac{1-\alpha_2}{2}\right)$	$\left(\frac{\alpha_2}{2}, 1 - \frac{\alpha_2}{2}\right)$

For $\alpha_1 \in (0, 1)$, no Nash equilibria in pure strategies exist since player *A*'s unique best response to $p_1 = i$ is $s_1 = 1 - i$ and player *B*'s unique best response to $s_1 = j$ is $p_1 = j$, for $i, j \in \{0, 1\}$. For $\alpha_1 \in \{0, 1\}$, the unique Nash equilibrium in pure strategies is given by $p = s = \alpha$. To find a mixed strategy equilibrium for any $\alpha \in [0, 1]$, we first determine a strategy p of player *B* that makes player *A* indifferent between her two pure strategies $s_1 = 1$ and $s_1 = 0$. Likewise, we determine the strategy s of player *A* that makes player *B* indifferent between his two pure strategies.

$$\begin{aligned}
 u_A(1, p_1) \stackrel{!}{=} u_A(0, p_1) &\Leftrightarrow \frac{1}{2} + \frac{\alpha_1 - p_1}{2} = \frac{1}{2} + \frac{\alpha_2 - p_2}{2} \\
 &\Leftrightarrow \alpha_1 - p_1 = \alpha_2 - p_2 \Leftrightarrow \alpha_1 = p_1 \\
 u_B(s_1, 1) \stackrel{!}{=} u_B(s_1, 0) &\Leftrightarrow \frac{1}{2} + \frac{1}{2}(\alpha_1 - 1 + 2s_1(1 - \alpha_1)) = \frac{1}{2} + \frac{1}{2}(1 - 2s_1)\alpha_1 \\
 &\Leftrightarrow 2s_1 - 1 = 0 \Leftrightarrow s_1 = \frac{1}{2}
 \end{aligned}$$

As $s_2 = 1 - s_1$ and $\alpha_2 = 1 - \alpha_1$, this yields $s = (\frac{1}{2}, \frac{1}{2})$ and $p = \alpha$. This completes the proof. □

Derivation of formulas (2.4) and (2.5) - Computation of the replicator dynamics.

Using the payoffs given in Table 2.1, we compute the replicator dynamics for s_1 and p_1 :

$$\begin{aligned}
\dot{s}_1 &= s_1(u_A(1, p_1) - u_A(s_1, p_1)) = s_1 \left(p_1 \frac{\alpha_1}{2} + (1 - p_1) \frac{1 + \alpha_1}{2} \right. \\
&\quad \left. - s_1 p_1 \frac{\alpha_1}{2} - s_1(1 - p_1) \frac{1 + \alpha_1}{2} - (1 - s_1) p_1 \frac{1 + \alpha_2}{2} - (1 - s_1)(1 - p_1) \frac{\alpha_2}{2} \right) \\
&= s_1(1 - s_1) \left(p_1 \frac{\alpha_1}{2} + (1 - p_1) \frac{1 + \alpha_1}{2} - p_1(1 - \frac{\alpha_1}{2}) - (1 - p_1) \frac{1 - \alpha_1}{2} \right) \\
&= s_1(1 - s_1) (-p_1 + \alpha_1) = s_1(1 - s_1) (\alpha_1 - p_1) \\
\dot{p}_1 &= p_1(u_B(s_1, 1) - u_B(s_1, p_1)) = p_1 \left(s_1(1 - \frac{\alpha_1}{2}) + (1 - s_1) \frac{1 - \alpha_2}{2} - p_1 s_1(1 - \frac{\alpha_1}{2}) \right. \\
&\quad \left. - p_1(1 - s_1) \frac{1 - \alpha_2}{2} - (1 - p_1) s_1 \frac{1 - \alpha_1}{2} - (1 - p_1)(1 - s_1) (1 - \frac{\alpha_2}{2}) \right) \\
&= p_1(1 - p_1) \left(s_1(1 - \frac{\alpha_1}{2}) + (1 - s_1) \frac{\alpha_1}{2} - s_1 \frac{1 - \alpha_1}{2} - (1 - s_1) \frac{1 + \alpha_1}{2} \right) \\
&= p_1(1 - p_1) (s_1 - \frac{1}{2})
\end{aligned}$$

□

Proof of Theorem 1. We are using the eigenvalue technique as explained in Friedman and Sinervo (2016). To assess the stability of a steady state x^* , we mul-

tiple the projection matrix $P_0 = \frac{1}{3} \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}$ by the Jacobian matrix

$J(x^*) = ((\frac{\partial V_i(x)}{\partial x_j}))|_{x=x^*}$. The resulting matrix has at least one eigenvalue of zero and the remaining two eigenvalues are ordered by their real parts. The state x^* will then be

- (locally) asymptotically stable if the largest real part is negative,
- completely unstable if the smallest real part is positive
- a saddle point if one real part is negative and one is positive. If the largest real part is also zero, the stability of that state can not be assessed with this method.

The Jacobian matrix of a general state $x = (\alpha_1, s_1, p_1)$ is given by

$$J(x) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ s_1(1 - s_1) & \alpha_1 - p_1 - 2s_1(\alpha_1 - p_1) & -s_1(1 - s_1) \\ 0 & p_1(1 - p_1) & (s_1 - \frac{1}{2})(1 - 2p_1) \end{pmatrix}.$$

Plugging in, we compute that $J((0, 0, 0)) = J((1, 1, 1)) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} \end{pmatrix}$.

Consequently, we get that $P_0J((0, 0, 0)) = P_0J((1, 1, 1)) = \begin{pmatrix} \frac{2\kappa-2}{3} & \frac{-2\kappa+2}{3} & \frac{1}{6} \\ \frac{-\kappa+1}{3} & \frac{\kappa-1}{3} & \frac{1}{6} \\ \frac{-\kappa+1}{3} & \frac{\kappa-1}{3} & \frac{-1}{3} \end{pmatrix}$.

The eigenvalues of this matrix are given by $\lambda_1 = 0, \lambda_2 = \frac{-1}{3}$ and

$\lambda_3 = \kappa - 1 < 0$ (as κ is assumed to be smaller than 1). Hence, the two non-zero eigenvalues have a negative real part, which makes the states $(0, 0, 0)$ and $(1, 1, 1)$ asymptotically stable steady states. For the two states $(0, 0, 1)$ and $(1, 1, 0)$,

we get $J((0, 0, 1)) = J((1, 1, 0)) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ 0 & -1 & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix}$ and $P_0J((0, 0, 1)) =$

$$P_0J((1, 1, 0)) = \begin{pmatrix} \frac{2\kappa-2}{3} & \frac{-2\kappa+3}{3} & \frac{-1}{6} \\ \frac{-\kappa+1}{3} & \frac{\kappa-3}{3} & \frac{-1}{6} \\ \frac{-\kappa+1}{3} & \frac{\kappa}{3} & \frac{1}{3} \end{pmatrix}.$$

Since we are only interested in the signs of the real parts of the eigenvalues of this matrix, we can multiply it by 6 to get rid of the fractions and compute the eigenvalues of the resulting matrix. In other words, we solve the characteristic equation associated to that matrix:

$$\left| \begin{pmatrix} 4\kappa - 4 - \lambda & -4\kappa + 6 & -1 \\ -2\kappa + 2 & 2\kappa - 6 - \lambda & -1 \\ -2\kappa + 2 & 2\kappa & 2 - \lambda \end{pmatrix} \right| \stackrel{!}{=} 0$$

$$\Leftrightarrow (4\kappa - 4 - \lambda) [(2\kappa - 6 - \lambda)(2 - \lambda) + 2\kappa] + (4\kappa - 6) [(-2\kappa + 2)(2 - \lambda) - 2\kappa + 2]$$

$$- [(-2\kappa + 2)2\kappa - (-2\kappa + 2)(2\kappa - 6)] = 0$$

$$\Leftrightarrow -\lambda^3 - 8\lambda^2 + 2\kappa\lambda^2 + 8\lambda - 2\lambda\kappa = 0$$

$$\Rightarrow \lambda = 0 \vee \lambda^2 + 8\lambda - 2\kappa\lambda - 8 + 2\kappa = 0$$

$$\Rightarrow \lambda = 0 \vee \lambda = \kappa - 4 \pm \sqrt{(4 - \kappa)^2 + 8 - 2\kappa}$$

Now, since $\kappa \in [0, 1)$, the square root is larger than $4 - \kappa$ which gives us exactly one positive and one negative (real) eigenvalue. Hence, the points $(0,0,1)$ and $(1,1,0)$ are saddle points of the dynamical system. Lastly, we want to assess the stability properties of the point $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. For this, we note that

$$J((\frac{1}{2}, \frac{1}{2}, \frac{1}{2})) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ \frac{1}{4} & 0 & \frac{-1}{4} \\ 0 & \frac{1}{4} & 0 \end{pmatrix} \text{ and } P_0 J((\frac{1}{2}, \frac{1}{2}, \frac{1}{2})) = \begin{pmatrix} \frac{8\kappa-9}{12} & \frac{-8\kappa+7}{12} & \frac{1}{12} \\ \frac{-2\kappa+3}{6} & \frac{4\kappa-5}{12} & \frac{-1}{6} \\ \frac{-4\kappa+3}{12} & \frac{2\kappa-1}{6} & \frac{1}{12} \end{pmatrix}.$$

Again, we can multiply this matrix by 12 and determine the eigenvalues of the resulting matrix:

$$\begin{aligned} & \left| \begin{pmatrix} 8\kappa - 9 - \lambda & -8\kappa + 7 & 1 \\ -4\kappa + 6 & 4\kappa - 5 - \lambda & -2 \\ -4\kappa + 3 & 4\kappa - 2 & 1 - \lambda \end{pmatrix} \right| \stackrel{!}{=} 0 \\ \Leftrightarrow & (8\kappa - 9 - \lambda)[(4\kappa - 5 - \lambda)(1 - \lambda) + 8\kappa - 4] + (8\kappa - 7)[(-4\kappa + 6)(1 - \lambda) - 8\kappa + 6] \\ & + [(-4\kappa + 6)(4\kappa - 2) - (-4\kappa + 3)(4\kappa - 5 - \lambda)] = 0 \\ \Leftrightarrow & -\lambda^3 + \lambda^2(12\kappa - 13) + \lambda(18 - 24\kappa) = 0 \\ \Leftrightarrow & \lambda = 0 \vee \underbrace{-\lambda^2 + \lambda(12\kappa - 13) + 18 - 24\kappa = 0}_{(I)} \end{aligned}$$

$$\begin{aligned} (I) & \Leftrightarrow \lambda^2 + \lambda(13 - 12\kappa) + 24\kappa - 18 = 0 \\ & \Rightarrow \lambda = 6\kappa - 6.5 \pm \sqrt{36\kappa^2 - 102\kappa + 60.25} \end{aligned}$$

This implies that for $36\kappa^2 - 102\kappa + 60.25 < 0$, the two non-zero eigenvalues are complex conjugates with real part $6\kappa - 6.5 < 0$. So in this case, the state will be asymptotically stable. In the other case, the state might be a saddle point if the real parts of the eigenvalues have different signs. The equation $36\kappa^2 - 102\kappa + 60.25 = 0$ is satisfied for $\kappa = \frac{51}{36} \pm \sqrt{\frac{1}{3}}$. Since the term is negative for $\kappa = 1$, it is negative for all $\kappa \in \left(\underbrace{\frac{51}{36} - \sqrt{\frac{1}{3}}}_{\approx 0.84}, 1 \right)$. For $\kappa \leq \frac{51}{36} - \sqrt{\frac{1}{3}}$, the two non-zero eigenvalues are real and one is negative for sure. So it remains to check when $6\kappa - 6.5 + \sqrt{36\kappa^2 - 102\kappa + 60.25} < 0$ holds (in which case the state will be asymptotically stable as well).

$$\begin{aligned}
& 6\kappa - 6.5 + \sqrt{36\kappa^2 - 102\kappa + 60.25} = 0 \\
& \Leftrightarrow 36\kappa^2 - 78\kappa + 42.25 = 36\kappa^2 - 102\kappa + 60.25 \\
& \Leftrightarrow 24\kappa = 18 \Leftrightarrow \kappa = \frac{3}{4}
\end{aligned}$$

This shows that for $\kappa > \frac{3}{4}$, both non-zero eigenvalues are negative (or have negative real parts) and the state $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is asymptotically stable. For $\kappa < \frac{3}{4}$ on the other hand, the two eigenvalues have different signs and the state $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is a saddle point. \square

Proof of Theorem 2. As in the proof of Theorem 1, our goal is to find the stationary points of the dynamical system and assess their stability properties afterwards. For this, it first makes sense to think about when the maxima and minima will attain which values depending on the relation of α and q . In the case of player A being the stronger player, we observe that

$$\begin{aligned}
& \min\left\{1, \frac{1}{1+q_B} + \frac{\alpha_1}{2}\right\} = 1 \Leftrightarrow \max\left\{0, \frac{q_B}{1+q_B} - \frac{\alpha_1}{2}\right\} = 0 \\
& \Leftrightarrow \alpha_1 \geq \frac{2q_B}{1+q_B} =: \tilde{q}, \text{ as well as} \\
& \min\left\{1, \frac{1}{1+q_B} + \frac{1-\alpha_1}{2}\right\} = 1 \Leftrightarrow \max\left\{0, \frac{q_B}{1+q_B} - \frac{1-\alpha_1}{2}\right\} = 0 \\
& \Leftrightarrow \alpha_1 \leq 1 - \tilde{q}.
\end{aligned}$$

Similarly, in the case of player B being the stronger player, we get the relations

$$\begin{aligned}
& \max\left\{0, \frac{q_A}{1+q_A} - \frac{\alpha_1}{2}\right\} = 0 \Leftrightarrow \min\left\{1, \frac{1}{1+q_A} + \frac{\alpha_1}{2}\right\} = 1 \\
& \Leftrightarrow \alpha_1 \geq \frac{2q_A}{1+q_A} =: \hat{q} \text{ and} \\
& \max\left\{0, \frac{q_A}{1+q_A} + \frac{\alpha_1 - 1}{2}\right\} = 0 \Leftrightarrow \min\left\{1, \frac{1}{1+q_A} + \frac{1-\alpha_1}{2}\right\} = 1 \\
& \Leftrightarrow \alpha_1 \leq 1 - \hat{q}.
\end{aligned}$$

Furthermore, it is important to note that $\tilde{q} = 1 - \tilde{q} \Leftrightarrow q_B = \frac{1}{3}$, hence $\tilde{q} < 1 - \tilde{q}$ for $q_B < \frac{1}{3}$ and $1 - \tilde{q} < \tilde{q}$ for $q_B > \frac{1}{3}$. Likewise, $\hat{q} < 1 - \hat{q}$ for $q_A < \frac{1}{3}$ and $1 - \hat{q} < \hat{q}$ for $q_A > \frac{1}{3}$. All in all, we will in the following distinguish between the two subcases 1. *Player A is the stronger player* and 2. *Player B is the stronger player*. Within

these cases, we will make further distinctions based on q_i and the relation of α_1 to \hat{q} and \tilde{q} . Before we start with this, we observe that for a general payoff matrix of the form

		Player B	
		$p_1 = 1$	$p_2 = 1$
Player A	$s_1 = 1$	(a_{11}, b_{11})	(a_{12}, b_{12})
	$s_2 = 1$	(a_{21}, b_{21})	(a_{22}, b_{22})

the replicator equation for s_1 reduces to

$$\begin{aligned} \dot{s}_1 &= s_1 (u_A(1, p_1) - u_A(s_1, p_1)) \\ &= s_1 (p_1 a_{11} + (1 - p_1) a_{12} - (s_1 p_1 a_{11} + s_1 (1 - p_1) a_{12} + (1 - s_1) p_1 a_{21} + (1 - s_1) (1 - p_1) a_{22})) \\ &= s_1 (1 - s_1) (p_1 a_{11} + (1 - p_1) a_{12} - p_1 a_{21} - (1 - p_1) a_{22}). \end{aligned}$$

Likewise, we get $\dot{p}_1 = p_1 (1 - p_1) (s_1 b_{11} + (1 - s_1) b_{21} - s_1 b_{12} - (1 - s_1) b_{22})$. Considering these formulas and recalling that $\dot{\alpha}_1 = (s_1 - \alpha_1)(1 - \kappa)$, we see that the points $(0, 0, 0)$, $(0, 0, 1)$, $(1, 1, 0)$ and $(1, 1, 1)$ will be stationary points of the dynamical system in all cases. In the following case distinction, we will thus be interested in finding additional interior stationary points (satisfying the necessary condition $s_1 = \alpha_1$). Subsequently, we will study the stability properties of the stationary points we found.

Case 1. i): Player A is stronger, $q_B < \frac{1}{3}$, $\alpha_1 \in [0, \tilde{q})$

In this case, the normal form game for a given experience vector α satisfying $\alpha_1 \in [0, \tilde{q})$ is given by

		$p_1 = 1$	$p_1 = 0$
		$(\frac{1}{1+q_B} - \frac{1}{2} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{1}{2} - \frac{\alpha_1}{2})$	$(\frac{1}{1+q_B} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} - \frac{\alpha_1}{2})$
$s_1 = 1$			
$s_1 = 0$	$(1, 0)$	$(\frac{1}{1+q_B} - \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{\alpha_1}{2})$	

This results in the following replicator equations:

$$\dot{s}_1 = s_1 (1 - s_1) \left(p_1 \left(\frac{1}{1 + q_B} - \frac{3}{2} - \frac{\alpha_1}{2} \right) + \alpha_1 \right) \quad (6)$$

$$\dot{p}_1 = p_1 (1 - p_1) \left(s_1 \left(\frac{1}{2} + \frac{q_B}{1 + q_B} + \frac{\alpha_1}{2} \right) - \frac{q_B}{1 + q_B} - \frac{\alpha_1}{2} \right) \quad (7)$$

Setting $\alpha_1 = s_1$ in (7) to find a potential interior steady state, we find that the big bracket is only zero for

$$s_1 = \alpha_1 = -\frac{q_B}{1+q_B} + \frac{\sqrt{3q_B^2 + 2q_B}}{1+q_B}$$

This term is zero for $q_B = 0$ and it is larger than \tilde{q} for $q_B \in (0, \frac{1}{3})$. Hence, the non-zero solution is not consistent with this case of $\alpha_1 \in [0, \tilde{q})$ and $\dot{p}_1 = 0$ only holds for $p_1 \in \{0, 1\}$. However, plugging in those values for p_1 in (6) shows that the big bracket is zero only for $\alpha_1 = 0$ and $p_1 = 0$ whereas it is always negative for $p_1 = 1$. Hence, there are no interior steady states in this case and we only have to assess the stability of the points $(0, 0, 0)$ and $(0, 0, 1)$. The Jacobian matrix of the point $(\alpha_1, s_1, p_1) = (0, 0, 0)$ is given by

$$J((0, 0, 0)) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{q_B}{1+q_B} \end{pmatrix}, \text{ yielding}$$

$$P_0 J((0, 0, 0)) = \frac{1}{3} \begin{pmatrix} 2\kappa - 2 & 2 - 2\kappa & \frac{q_B}{1+q_B} \\ 1 - \kappa & \kappa - 1 & \frac{q_B}{1+q_B} \\ 1 - \kappa & \kappa - 1 & -\frac{2q_B}{1+q_B} \end{pmatrix}.$$

Since we are only interested in the signs of the real parts of the eigenvalues of this matrix, we multiply it by 3 and solve the corresponding characteristic equation of the resulting matrix. This gives us the following solutions:

$$\lambda_1 = 0 \wedge \lambda_{2,3} = \frac{3\kappa - 3 - \frac{2q_B}{1+q_B}}{2} \pm \sqrt{\left(\frac{3\kappa - 3 - \frac{2q_B}{1+q_B}}{2}\right)^2 - \frac{6q_B - 6\kappa q_B}{1+q_B}}$$

Now, as $\kappa \in [0, 1)$, we see that $\frac{3\kappa - 3 - \frac{2q_B}{1+q_B}}{2} < 0$ and $\frac{6q_B - 6\kappa q_B}{1+q_B} > 0$. Consequently, we do not have to compute when the term under the root is positive or negative: When it is positive, the root will have an absolute value which is smaller than the absolute value of $\frac{3\kappa - 3 - \frac{2q_B}{1+q_B}}{2}$. When it is negative, the real parts of both eigenvalues λ_2 and λ_3 are given by $\frac{3\kappa - 3 - \frac{2q_B}{1+q_B}}{2}$. In total, both non-zero eigenvalues have negative real parts and the point $(0, 0, 0)$ is an asymptotically stable state of the dynamical

system. For the state $(0, 0, 1)$, we first compute the matrix $P_0J((0, 0, 1))$. This yields

$$P_0J((0, 0, 1)) = \frac{1}{3} \begin{pmatrix} 2\kappa - 2 & \frac{7}{2} - 2\kappa - \frac{1}{1+q_B} & \frac{-q_B}{1+q_B} \\ 1 - \kappa & \kappa - 4 + \frac{2}{1+q_B} & \frac{-q_B}{1+q_B} \\ 1 - \kappa & \kappa + \frac{1}{2} - \frac{1}{1+q_B} & \frac{2q_B}{1+q_B} \end{pmatrix}.$$

We can multiply this matrix by 3 and compute the eigenvalues of the resulting matrix (remember that we are only interested in the sign of the real parts of the eigenvalues). This leads to the eigenvalues $\lambda_1 = 0$ and

$$\lambda_{2,3} = \frac{3\kappa - 4}{2} \pm \sqrt{\left(\frac{3\kappa - 4}{2}\right)^2 + \frac{3(\kappa + q_B - 1 - (\kappa - 4)q_B^2)}{2(q_B + 1)^2}}.$$

The real parts of the two non-zero eigenvalues are negative if and only if $\kappa < \frac{1 - q_B - 4q_B^2}{1 - q_B^2}$, making the point $(0, 0, 1)$ stable in exactly this case.

Case 1. ii): Player A is stronger, $q_B < \frac{1}{3}$, $\alpha_1 \in [\tilde{q}, 1 - \tilde{q}]$

In this case, the game looks like

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$\left(\frac{1}{1+q_B} - \frac{1}{2} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{1}{2} - \frac{\alpha_1}{2}\right)$	$(1, 0)$
$s_1 = 0$	$(1, 0)$	$\left(\frac{1}{1+q_B} - \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{\alpha_1}{2}\right)$

and we get

$$\dot{s}_1 = s_1(1 - s_1) \left(p_1 \left(\frac{2}{1+q_B} - \frac{5}{2} \right) + 1 - \frac{1}{1+q_B} + \frac{\alpha_1}{2} \right) \quad (8)$$

$$\dot{p}_1 = p_1(1 - p_1) \left(s_1 \left(\frac{2q_B}{1+q_B} + \frac{1}{2} \right) - \frac{q_B}{1+q_B} - \frac{\alpha_1}{2} \right). \quad (9)$$

The unique interior solution (meaning that both p_1 and s_1 are not in $\{0, 1\}$) of this system is $(\alpha_1, s_1, p_1) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Hence, we only need to check the stability properties of this point. Performing the usual computations yields

$$P_0J\left(\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)\right) = \frac{1}{3} \begin{pmatrix} 2\kappa - 2 & \frac{15}{8} - 2\kappa - \frac{q_B}{2+2q_B} & \frac{5}{8} - \frac{1}{2+2q_B} \\ \frac{11}{8} - \kappa & \kappa - \frac{9}{8} - \frac{q_B}{2+2q_B} & \frac{1}{1+q_B} - \frac{10}{8} \\ \frac{5}{8} - \kappa & \kappa - \frac{3}{4} + \frac{q_B}{1+q_B} & \frac{5}{8} - \frac{1}{2+2q_B} \end{pmatrix}.$$

The non-zero eigenvalues of the product of 3 and this matrix are given by

$$\lambda_{2,3} = \frac{3\kappa - 3}{2} \pm \sqrt{\left(\frac{3\kappa - 3}{2}\right)^2 + \frac{3(21q_B^2 + 58q_B + 21)}{64(q_B + 1)^2} - \frac{3\kappa(7q_B + 3)}{8(q_B + 1)}}.$$

These two values both have negative real parts (meaning that $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ is stable) if and only if $\kappa > \frac{21q_B^2 + 58q_B + 21}{8(7q_B + 3)(q_B + 1)}$

holds.

Case 1. iii): Player A is stronger, $q_B < \frac{1}{3}$, $\alpha_1 \in (1 - \tilde{q}, 1]$

The game is symmetric with respect to the two options, so this is equal to case 1. i).

Case 1. iv): Player A is stronger, $q_B > \frac{1}{3}$, $\alpha_1 \in [0, 1 - \tilde{q}]$ In this case, the game looks exactly like the one in case 1. i). This is logical since in both cases, $\alpha_1 < \tilde{q}$ and $\alpha_1 < 1 - \tilde{q}$ hold and only the order of \tilde{q} and $1 - \tilde{q}$ is reversed. Consequently, the replicator equations and the steady states are the same, as well as the stability properties of these steady states (since those properties are derived from the replicator equations as well via the Jacobian matrix and the eigenvalues of its product with a projection matrix, and nothing in the computations hinged on the assumption $q_B < \frac{1}{3}$).

Case 1. v): Player A is stronger, $q_B > \frac{1}{3}$, $\alpha_1 \in (1 - \tilde{q}, \tilde{q})$

In this case, the normal form game for a given experience vector α satisfying $\alpha_1 \in (1 - \tilde{q}, \tilde{q})$ is given by

$$\begin{array}{r}
 \begin{array}{c} p_1 = 1 \\ p_1 = 0 \end{array}
 \begin{array}{|c|c|}
 \hline
 & p_1 = 0 \\
 \hline
 s_1 = 1 & \left(\frac{1}{1+q_B} - \frac{1}{2} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{1}{2} - \frac{\alpha_1}{2} \right) \\
 s_1 = 0 & \left(\frac{1}{1+q_B} + \frac{1-\alpha_1}{2}, \frac{q_B}{1+q_B} - \frac{1-\alpha_1}{2} \right) \\
 \hline
 & \left(\frac{1}{1+q_B} + \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} - \frac{\alpha_1}{2} \right) \\
 & \left(\frac{1}{1+q_B} - \frac{\alpha_1}{2}, \frac{q_B}{1+q_B} + \frac{\alpha_1}{2} \right) \\
 \hline
 \end{array}
 \end{array}$$

This results in the following replicator equations:

$$\begin{aligned}
 \dot{s}_1 &= s_1(1 - s_1)(\alpha_1 - p_1) \\
 \dot{p}_1 &= p_1(1 - p_1)\left(s_1 - \frac{1}{2}\right)
 \end{aligned}$$

We can see that these equations are exactly the same as (2.4) and (2.5). In particular, they do not depend on the new parameter q_B (which is logical, since only the baseline probability of winning depends on q_B , but this is constant for each player among the four cells in the normal form matrix). Consequently, the only interior steady state is given by $(\alpha_1, s_1, p_1) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ and the result of the stability analysis of this point is the same as in the proof of Theorem 1: It is asymptotically stable for $\kappa > \frac{3}{4}$ and a saddle point for $\kappa < \frac{3}{4}$.

Case 1. vi): Player A is stronger, $q_B > \frac{1}{3}$, $\alpha_1 \in [\tilde{q}, 1]$

Due to symmetry, this case is the same as 1. iv).

Case 2. i): Player B is stronger, $q_A < \frac{1}{3}$, $\alpha_1 \in [0, \hat{q}]$

In this case, the game looks like

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(0, 1)$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1}{2}, \frac{1}{1+q_A} - \frac{\alpha_1}{2})$
$s_1 = 0$	$(\frac{q_A}{1+q_A} + \frac{1-\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1-1}{2})$	$(\frac{q_A}{1+q_A} - \frac{\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1}{2})$

and we get

$$\dot{s}_1 = s_1(1 - s_1) \left(p_1 \left(-\frac{1}{2} - \frac{\alpha_1}{2} - \frac{q_A}{1+q_A} \right) + \alpha_1 \right) \quad (10)$$

$$\dot{p}_1 = p_1(1 - p_1) \left(s_1 \left(1 - \frac{1}{1+q_A} + \frac{\alpha_1 + 1}{2} \right) - \frac{1}{2} \right). \quad (11)$$

Recall that in an interior solution, we will have $s_1 = \alpha_1$ due to (2.3). Using this and the requirement $s_1 \in [0, 1]$ in equation (11), we get that

$$s_1 = \frac{1}{1+q_A} - \frac{3}{2} + \sqrt{\left(\frac{1}{1+q_A} - \frac{3}{2} \right)^2 + 1}$$

is the unique solution candidate for an interior steady state. However, this solution candidate is strictly larger than \hat{q} for $q_A < \frac{1}{3}$, which would contradict $\alpha_1 \in [0, \hat{q}]$. So to get $\dot{p}_1 = 0$, we need $p_1 = 0$ or $p_1 = 1$ in (11). Turning to (10), we see that for $p_1 = 0$, the big bracket is zero only for $\alpha_1 = 0$ while for $p_1 = 1$, the big bracket can never be zero. Hence, there is no interior steady state in this case and we only need to check the stability of the points $(0, 0, 0)$ and $(0, 0, 1)$. The corresponding points $(1, 1, 1)$ and $(1, 1, 0)$ will then have the same stability properties due to symmetry. Computing the Jacobian matrix associated to the point $(0, 0, 0)$ yields

$$J((0, 0, 0)) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} \end{pmatrix}.$$

Since this is the same Jacobian matrix we obtained in the proof of Theorem 1, we can conclude that the point $(0, 0, 0)$ is an asymptotically stable state in this case as well. The Jacobian matrix associated to the point $(0, 0, 1)$ is given by

$$J((0, 0, 1)) = \begin{pmatrix} \kappa - 1 & 1 - \kappa & 0 \\ 0 & -\frac{1}{2} - \frac{q_A}{1+q_A} & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix}$$

and the eigenvalues of $3 * P_0J((0, 0, 1))$ are given by $\lambda_0 = 0$ and

$$\lambda_{2,3} = \frac{3\kappa(q_A + 1) - 5q_A - 3}{2q_A + 2} \pm \sqrt{\left(\frac{3\kappa(q_A + 1) - 5q_A - 3}{2q_A + 2}\right)^2 + \frac{3(2\kappa(q_A - 1) + q_A + 3)}{4(q_A + 1)}}.$$

As $\frac{3\kappa(q_A+1)-5q_A-3}{2q_A+2} < 0$ and $\frac{3(2\kappa(q_A-1)+q_A+3)}{4(q_A+1)} > 0$, exactly one of the two non-zero eigenvalues is positive and one is negative. Consequently, the state $(0, 0, 1)$ is a saddle point of the dynamical system.

Case 2. ii): Player B is stronger, $q_A < \frac{1}{3}$, $\alpha_1 \in [\hat{q}, 1 - \hat{q}]$

In this case, the game looks like

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(0, 1)$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1}{2}, \frac{1}{1+q_A} - \frac{\alpha_1}{2})$
$s_1 = 0$	$(\frac{q_A}{1+q_A} + \frac{1-\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1-1}{2})$	$(0, 1)$

and we get

$$\dot{s}_1 = s_1(1 - s_1) \left(\frac{q_A}{1 + q_A} + \frac{\alpha_1}{2} - p_1 \left(\frac{2q_A}{1 + q_A} + \frac{1}{2} \right) \right) \quad (12)$$

$$\dot{p}_1 = p_1(1 - p_1) \left(s_1 \left(\frac{5}{2} - \frac{2}{1 + q_A} \right) - \frac{q_A}{1 + q_A} + \frac{\alpha_1 - 1}{2} \right). \quad (13)$$

To find an interior steady state, we plug in $\alpha_1 = s_1$ into (13) and obtain $s_1 = \frac{1}{2}$ as the unique solution candidate. Plugging in $s_1 = \alpha_1 = \frac{1}{2}$ into (12) yields $p_1 = \frac{1}{2}$ as well, making $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ the unique interior steady state in this case. To assess its stability, we perform the usual computations and get

$$P_0J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right) = \frac{1}{3} \begin{pmatrix} 2\kappa - \frac{9}{4} & \frac{11}{8} - 2\kappa - \frac{1}{2+2q_A} & \frac{1}{8} + \frac{q_A}{2+2q_A} \\ \frac{9}{8} - \kappa & \kappa - \frac{13}{8} + \frac{1}{2+2q_A} & \frac{-q_A}{1+q_A} - \frac{1}{4} \\ \frac{9}{8} - \kappa & \kappa + \frac{1}{4} - \frac{1}{1+q_A} & \frac{1}{8} - \frac{q_A}{2+2q_A} \end{pmatrix}.$$

The non-zero eigenvalues of the product of 3 and this matrix are given by

$$\lambda_{2,3} = \frac{3\kappa - \frac{13}{4}}{2} \pm \sqrt{\left(\frac{3\kappa - \frac{13}{4}}{2}\right)^2 + \frac{3(3q_A+7)(5q_A+1)}{64(q_A+1)^2} + \kappa\left(\frac{3}{2(q_A+1)} - \frac{15}{8}\right)}. \text{ These two values both have negative real parts (meaning that } (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) \text{ is stable) if and only if } \kappa > \frac{21q_A^2+58q_A+21}{8(7q_A+3)(q_A+1)} \text{ holds.}$$

Case 2. iii): Player B is stronger, $q_A < \frac{1}{3}$, $\alpha_1 \in (1 - \hat{q}, 1]$

Due to symmetry, this case is the same as 2. i).

Case 2. iv): Player B is stronger, $q_A > \frac{1}{3}$, $\alpha_1 \in [0, 1 - \hat{q}]$

In this case, the game looks like

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(0, 1)$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1}{2}, \frac{1}{1+q_A} - \frac{\alpha_1}{2})$
$s_1 = 0$	$(\frac{q_A}{1+q_A} + \frac{1-\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1-1}{2})$	$(\frac{q_A}{1+q_A} - \frac{\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1}{2})$

and we get

$$\dot{s}_1 = s_1(1 - s_1) \left(p_1 \left(-\frac{1}{2} - \frac{\alpha_1}{2} - \frac{q_A}{1+q_A} \right) + \alpha_1 \right) \quad (14)$$

$$\dot{p}_1 = p_1(1 - p_1) \left(s_1 \left(1 - \frac{1}{1+q_A} + \frac{\alpha_1 + 1}{2} \right) - \frac{1}{2} \right). \quad (15)$$

Recall that in an interior solution, we will have $s_1 = \alpha_1$ due to (2.3). Using this and the requirement $s_1 \in [0, 1]$ in equation (15), we get that

$$s_1 = \frac{1}{1+q_A} - \frac{3}{2} + \sqrt{\left(\frac{1}{1+q_A} - \frac{3}{2} \right)^2 + 1}$$

is the unique solution candidate for an interior steady state. However, this solution candidate is strictly larger than $1 - \hat{q}$ for $q_A > \frac{1}{3}$, which would contradict $\alpha_1 \in [0, 1 - \hat{q}]$. So to get $\dot{p}_1 = 0$, we need $p_1 = 0$ or $p_1 = 1$ in (15). Turning to (14), we see that for $p_1 = 0$, the big bracket is zero only for $\alpha_1 = 0$ while for $p_1 = 1$, the big bracket can never be zero. Hence, there is no interior steady state in this case and we only need to check the stability of the points $(0, 0, 0)$ and $(0, 0, 1)$. The corresponding points $(1, 1, 1)$ and $(1, 1, 0)$ will then have the same stability properties due to symmetry.

We can see that the matrix and the replicator equations are the same as in case 2. i). As the assumption $q_A < \frac{1}{3}$ was not needed in the stability analysis in 2. i), we can conclude that the point $(0, 0, 0)$ is asymptotically stable and the point $(0, 0, 1)$ is a saddle point, exactly as in case 2. i).

Case 2. v): Player B is stronger, $q_A > \frac{1}{3}$, $\alpha_1 \in (1 - \hat{q}, \hat{q})$

In this case, the game looks like

	$p_1 = 1$	$p_1 = 0$
$s_1 = 1$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1-1}{2}, \frac{1}{1+q_A} + \frac{1-\alpha_1}{2})$	$(\frac{q_A}{1+q_A} + \frac{\alpha_1}{2}, \frac{1}{1+q_A} - \frac{\alpha_1}{2})$
$s_1 = 0$	$(\frac{q_A}{1+q_A} + \frac{1-\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1-1}{2})$	$(\frac{q_A}{1+q_A} - \frac{\alpha_1}{2}, \frac{1}{1+q_A} + \frac{\alpha_1}{2})$

and we get

$$\begin{aligned}\dot{s}_1 &= s_1(1 - s_1)(\alpha_1 - p_1) \\ \dot{p}_1 &= p_1(1 - p_1)\left(s_1 - \frac{1}{2}\right).\end{aligned}$$

Hence, the only interior steady state is $(\alpha_1, s_1, p_1) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ and its stability properties are as in case 1. v).

Case 2. vi): Player B is stronger, $q_A > \frac{1}{3}$, $\alpha_1 \in [\hat{q}, 1]$

Due to symmetry, this case is the same as 2. iv). \square

Proof of Lemma 2. Using the payoffs given in Table 2.4, we can compute the replicator equations for s_1, s_2, p_1 and p_2 . The steady states of the system are exactly those points where the growth rates of all variables are zero. Concretely, we have to find the points $(\alpha_1, \alpha_2, s_1, s_2, p_1, p_2)$ that simultaneously satisfy the following conditions (16) to (21):

$$\dot{\alpha}_1 = (s_1 - \alpha_1)(1 - \kappa) \stackrel{!}{=} 0 \quad (16)$$

$$\dot{\alpha}_2 = (s_2 - \alpha_2)(1 - \kappa) \stackrel{!}{=} 0 \quad (17)$$

$$\dot{s}_1 = s_1 \left((1 - s_1) \frac{2\alpha_1 + \alpha_2 - 2p_1 - p_2}{2} - s_2 \frac{2\alpha_2 + \alpha_1 - 2p_2 - p_1}{2} \right) \stackrel{!}{=} 0 \quad (18)$$

$$\dot{s}_2 = s_2 \left((1 - s_2) \frac{\alpha_1 + 2\alpha_2 - p_1 - 2p_2}{2} - s_1 \frac{2\alpha_1 + \alpha_2 - 2p_1 - p_2}{2} \right) \stackrel{!}{=} 0 \quad (19)$$

$$\dot{p}_1 = p_1 \left((1 - p_1) \left(s_1 + \frac{s_2 - 1}{2} \right) - p_2 \left(s_2 + \frac{s_1 - 1}{2} \right) \right) \stackrel{!}{=} 0 \quad (20)$$

$$\dot{p}_2 = p_2 \left((1 - p_2) \left(s_2 + \frac{s_1 - 1}{2} \right) - p_1 \left(s_1 + \frac{s_2 - 1}{2} \right) \right) \stackrel{!}{=} 0 \quad (21)$$

The equations (16) and (17) are satisfied if and only if $s_1 = \alpha_1$ and $s_2 = \alpha_2$ hold. From there, one can go through all the cases where one of the remaining equations holds and check in which cases all the other growth rates are zero as well. For instance, the relation $\dot{p}_1 = 0$ (equation (20)) holds in the following cases:

- i) $p_1 = 0$
- ii) $p_1 = 1$
- iii) $p_1 \notin \{0, 1\}, p_2 = 0, s_1 = \frac{1-s_2}{2}$

$$\text{iv) } p_1 \notin \{0, 1\}, p_2 \neq 0, s_1 = s_2 = s_3 = \frac{1}{3}$$

$$\text{v) } (1 - p_1)(s_1 + \frac{s_2 - 1}{2}) = p_2(s_2 + \frac{s_1 - 1}{2}), \text{ all factors } \neq 0$$

Going through each of the above cases while satisfying equations (18), (19) and (21) as well yields the set of steady states given in Lemma 2. \square

Proof of Theorem 3. We apply the same technique as in the proof of Theorem 1. First, the Jacobian matrices at a steady state are calculated and this time multiplied (from the left) by the six-dimensional projection matrix

$$P_0 = \frac{1}{6} \begin{pmatrix} 5 & -1 & -1 & -1 & -1 & -1 \\ -1 & 5 & -1 & -1 & -1 & -1 \\ -1 & -1 & 5 & -1 & -1 & -1 \\ -1 & -1 & -1 & 5 & -1 & -1 \\ -1 & -1 & -1 & -1 & 5 & -1 \\ -1 & -1 & -1 & -1 & -1 & 5 \end{pmatrix}.$$

Thereafter, we compute the real parts of the resulting matrices' eigenvalues. Due to the projection, one eigenvalue will always be zero. The steady state in question is asymptotically stable whenever the second highest real part of the eigenvalues is negative. For the categories i), ii), iii) and v), it is sufficient to make the computations for one representative, since all steady states in these categories only differ by relabeling the options. For the representative of category iv), we have to deal with all the possible values $p_2 \in [0, 1]$. The Jacobian matrix $J(x)$ for a state $x = (\alpha_1, \alpha_2, s_1, s_2, p_1, p_2)$ is given by

$$\left(\begin{array}{cccccc}
\kappa-1 & 0 & 1-\kappa & 0 & 0 & 0 \\
0 & \kappa-1 & 0 & 1-\kappa & 0 & 0 \\
s_1(1-s_1) & \frac{s_1(1-s_1)}{2} & (1-2s_1)\frac{2\alpha_1+\alpha_2-2p_1-p_2}{2} & -S_1\frac{2\alpha_2+\alpha_1-2p_2-p_1}{2} & -s_1(1-s_1) & -\frac{s_1(1-s_1)}{2} \\
-\frac{s_1s_2}{2} & -s_1s_2 & -s_2\frac{2\alpha_2+\alpha_1-2p_2-p_1}{2} & & +\frac{s_1s_2}{2} & +s_1s_2 \\
\frac{s_2(1-s_2)}{2} & s_2(1-s_2) & -S_2\frac{2\alpha_1+\alpha_2-2p_1-p_2}{2} & (1-2s_2)\frac{2\alpha_2+\alpha_1-2p_2-p_1}{2} & -\frac{s_2(1-s_2)}{2} & -s_2(1-s_2) \\
-s_1s_2 & -\frac{s_1s_2}{2} & & -s_1\frac{2\alpha_1+\alpha_2-2p_1-p_2}{2} & +s_1s_2 & +\frac{s_1s_2}{2} \\
0 & 0 & p_1(1-p_1)-\frac{p_1p_2}{2} & \frac{p_1(1-p_1)}{2}-p_1p_2 & (1-2p_1)(s_1+\frac{s_2-1}{2}) & -p_1(s_2+\frac{s_1-1}{2}) \\
& & & & -p_2(\frac{s_1-1}{2}+s_2) & \\
0 & 0 & \frac{p_2(1-p_2)}{2}-p_1p_2 & p_2(1-p_2)-\frac{p_1p_2}{2} & -p_2(s_1+\frac{s_2-1}{2}) & (1-2p_2)(s_2+\frac{s_1-1}{2}) \\
& & & & & -p_1(\frac{s_2-1}{2}+s_1)
\end{array} \right)$$


Concretely, we get the following matrices for the representatives of the categories i), ii), iii) and v):

$$P_0 * J\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) = \frac{1}{6} \begin{pmatrix}
5\kappa - \frac{31}{6} & \frac{5}{6} - \kappa & \frac{29}{6} - 5\kappa & \kappa - \frac{7}{6} & \frac{1}{6} & \frac{1}{6} \\
\frac{5}{6} - \kappa & 5\kappa - \frac{31}{6} & \kappa - \frac{7}{6} & \frac{29}{6} - 5\kappa & \frac{1}{6} & \frac{1}{6} \\
\frac{11}{6} - \kappa & \frac{5}{6} - \kappa & \kappa - \frac{7}{6} & \kappa - \frac{7}{6} & -\frac{5}{6} & \frac{1}{6} \\
\frac{5}{6} - \kappa & \frac{11}{6} - \kappa & \kappa - \frac{7}{6} & \kappa - \frac{7}{6} & \frac{1}{6} & -\frac{5}{6} \\
\frac{5}{6} - \kappa & \frac{5}{6} - \kappa & \kappa - \frac{1}{6} & \kappa - \frac{7}{6} & \frac{1}{6} & \frac{1}{6} \\
\frac{5}{6} - \kappa & \frac{5}{6} - \kappa & \kappa - \frac{7}{6} & \kappa - \frac{1}{6} & \frac{1}{6} & \frac{1}{6}
\end{pmatrix},$$


$$\begin{aligned}
P_0 * J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0\right) &= \frac{1}{6} \begin{pmatrix} 5\kappa - 5 & 1 - \kappa & \frac{23}{4} - 5\kappa & \kappa - \frac{1}{4} & -\frac{1}{4} & -\frac{1}{4} \\ 1 - \kappa & 5\kappa - 5 & \kappa - \frac{1}{4} & \frac{23}{4} - 5\kappa & -\frac{1}{4} & -\frac{1}{4} \\ \frac{7}{4} - \kappa & \frac{1}{4} - \kappa & \kappa - \frac{5}{2} & \kappa - \frac{5}{2} & -1 & \frac{1}{2} \\ \frac{1}{4} - \kappa & \frac{7}{4} - \kappa & \kappa - \frac{5}{2} & \kappa - \frac{5}{2} & \frac{1}{2} & -1 \\ 1 - \kappa & 1 - \kappa & \kappa + \frac{1}{2} & \kappa - 1 & \frac{5}{4} & -\frac{1}{4} \\ 1 - \kappa & 1 - \kappa & \kappa - 1 & \kappa + \frac{1}{2} & -\frac{1}{4} & \frac{5}{4} \end{pmatrix}, \\
P_0 * J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right) &= \frac{1}{6} \begin{pmatrix} 5\kappa - 5 & 1 - \kappa & 5 - 5\kappa & \kappa - 1 & \frac{1}{4} & \frac{1}{4} \\ 1 - \kappa & 5\kappa - 5 & \kappa - 1 & 5 - 5\kappa & \frac{1}{4} & \frac{1}{4} \\ \frac{7}{4} - \kappa & \frac{1}{4} - \kappa & \kappa - 1 & \kappa - 1 & -\frac{1}{2} & 1 \\ \frac{1}{4} - \kappa & \frac{7}{4} - \kappa & \kappa - 1 & \kappa - 1 & 1 & -\frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - \frac{1}{4} & \kappa - \frac{7}{4} & -\frac{1}{2} & -\frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - \frac{7}{4} & \kappa - \frac{1}{4} & -\frac{1}{2} & -\frac{1}{2} \end{pmatrix}, \\
P_0 * J(1, 0, 1, 0, 1, 0) &= \frac{1}{6} \begin{pmatrix} 5\kappa - 5 & 1 - \kappa & 5 - 5\kappa & \kappa - 1 & \frac{1}{2} & \frac{1}{2} \\ 1 - \kappa & 5\kappa - 5 & \kappa - 1 & 5 - 5\kappa & \frac{1}{2} & \frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - 1 & \kappa - 1 & \frac{1}{2} & \frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - 1 & \kappa - 1 & \frac{1}{2} & \frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - 1 & \kappa - 1 & -\frac{5}{2} & \frac{1}{2} \\ 1 - \kappa & 1 - \kappa & \kappa - 1 & \kappa - 1 & \frac{1}{2} & -\frac{5}{2} \end{pmatrix}.
\end{aligned}$$

These matrices all depend on the parameter κ . To show that the states in category v) are asymptotically stable, we need to verify that the largest non-zero real part of the associated matrix's eigenvalues is negative for all $\kappa < 1$. Likewise, we need to verify that the largest non-zero real part of the eigenvalues belonging to the matrices associated to categories i) - iii) is positive for each category and all $\kappa < 1$. Since we are only interested in the sign of the largest real part of the eigenvalues, it is sufficient to compute the eigenvalues of each above matrix multiplied by 6 (this way, we will get rid of the factor $\frac{1}{6}$ in all cases). To do this, we compute the real parts of the eigenvalues of each of these matrices numerically, depending on the parameter κ . The plots are given in Figure 4.


For categories i) and iii), we see that the largest real part of the eigenvalues is positive and approaches zero for κ being close to 1. Calculating the eigenvalues of the matrices $6 * P_0 * J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$ and $6 * P_0 * J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0\right)$ for $\kappa = 1$ confirms these observations: The real parts of the eigenvalues of $6 * P_0 * J\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$ are




(a) Mixing between all 3 options, consistent preparation (category i))



(b) Mixing between 2 options, completely inconsistent preparation (category ii))



(c) Mixing between 2 options, consistent preparation (category iii))



(d) Perfect specialization on one option, consistent preparation (category v))

Figure 4: Plots of the real parts of the eigenvalues belonging to the matrices associated to categories i), ii), iii) and v)

$0, 0, 0, 0, 0, -1$, while the real parts of the eigenvalues of $6 * P_0 * J(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0)$ are $0, 0, 0, 0, \frac{1}{6}, \frac{1}{6}$. This implies that in both cases, the largest real part of all the eigenvalues is strictly positive for all values $\kappa \in [0, 1)$ and the states in categories i) and iii) are indeed saddle points. The eigenvalues of $6 * P_0 * J(1, 0, 1, 0, 1, 0)$ are $0, 0, 0, 0, -2$ and -3 . This implies that the largest non-zero eigenvalue is negative (for this matrix, all eigenvalues are real numbers) for all $\kappa \in [0, 1)$ since the absolute value of the second highest eigenvalue is decreasing and only hits zero at $\kappa = 1$. In conclusion, the states in category v) are asymptotically stable. Lastly, as the largest real value of the eigenvalues of the matrix associated to category ii) is always positive, the states in this category are saddle points.

For a general representative of category iv), the associated matrix is given by

$$P_0 J(1, 0, 1, 0, 0, p_2) = \frac{1}{6} \begin{pmatrix} 5\kappa - 5 & 1 - \kappa & 5 - 5\kappa + \frac{2-2p_2+p_2^2}{2} & \kappa - \frac{3p_2}{2} + p_2^2 & \frac{p_2-1}{2} & 0 \\ 1 - \kappa & 5\kappa - 5 & \kappa - p_2 + \frac{p_2^2}{2} & 6 - 5\kappa - \frac{3p_2}{2} + p_2^2 & \frac{p_2-1}{2} & 0 \\ 1 - \kappa & 1 - \kappa & \kappa - 6 + 2p_2 + \frac{p_2^2}{2} & \kappa - 3 + \frac{9p_2}{2} + p_2^2 & \frac{p_2-1}{2} & 0 \\ 1 - \kappa & 1 - \kappa & \kappa - p_2 + \frac{p_2^2}{2} & \kappa - 3 - \frac{9p_2}{2} + p_2^2 & \frac{p_2-1}{2} & 0 \\ 1 - \kappa & 1 - \kappa & \kappa - p_2 + \frac{p_2^2}{2} & \kappa - \frac{3p_2}{2} + p_2^2 & \frac{p_2+5}{2} & 0 \\ 1 - \kappa & 1 - \kappa & \kappa + \frac{4p_2-5p_2^2}{2} & \kappa - \frac{3p_2}{2} + p_2^2 & \frac{-5p_2-1}{2} & 0 \end{pmatrix}.$$

Our goal is again to compute the largest real part of the eigenvalues of $6 * P_0 J(1, 0, 1, 0, 0, p_2)$, depending on κ and p_2 . To achieve this, we first fix a value of p_2 and compute the maximal real part of all the eigenvalues depending on κ . Thereafter, we take the minimal value of all these maximal real parts over all values $\kappa \in [0, 1]$. This procedure is then repeated for every value $p_2 \in [0, 1]$. The result is plotted in Figure 5. The Python code used to generate Figures 4 and 5 can be found under <https://github.com/mgramb/evolutionary-preparation.git>.




Figure 5: For each p_2 , the plotted value is the minimum over all largest real parts of the eigenvalues for $\kappa \in [0, 1]$.

As the real parts of these eigenvalues are strictly positive for every p_2 , we can conclude that this holds for every combination of values $\kappa \in [0, 1)$ and $p_2 \in [0, 1]$. Similar to the above procedure, it is shown that the smallest real part of the

eigenvalues is negative for all combinations of $\kappa \in [0, 1)$ and $p_2 \in [0, 1]$. This means that the states in category iv) are always saddle points. This concludes the proof. \square

B Elements of Evolutionary Game Theory

In this section, we review several definitions and results from evolutionary game theory. A more comprehensive account of this field, with many applications in economics, can be found in Sandholm (2010) and Friedman and Sinervo (2016).

B.1 Population games

Population games are games in which individuals from different subpopulations play against each other. Traditionally, each individual in a subpopulation is programmed to play a particular pure strategy in the sense that they always play that strategy in the game. In the context of standard game theory, you can think of a normal form game where each player represents a different subpopulation and where opposing players are drawn randomly from each subpopulation. Let n be the number of subpopulations or the number of players in the normal form game. Also, let S_i be the number of pure strategies available to subpopulation $i \in \{1, \dots, n\}$. Then we can denote by x_{ih} the share of individuals from population i that are programmed to use pure strategy $h \in \{1, \dots, S_i\}$. Consequently, when a player faces another player from subpopulation i in the game, he will expect that player to play mixed strategy $x_i = (x_{i1}, \dots, x_{iS_i})$. Note that this implies that mathematically, it is not important whether each individual is assumed to play a pure or a mixed strategy, since the expected strategy of a random opponent from that population will still be a mixed strategy. The reason we assume each player to play a pure strategy in traditional evolutionary game theory is that this makes the most sense in many of the applications in biology and animal conflict. The vector x_i can be interpreted as the state of subpopulation i since it indicates the proportion of individuals playing each of the available pure strategies. This means that the state of the whole population can be described by the vector $x = (x_1, \dots, x_n)$, where x_i is the state of subpopulation i .

In these population games, the focus of study is on how the shares of individuals playing a particular strategy evolve in a repeated game, depending on how that strategy performs compared to other strategies in the population. These evolutions

can be captured by growth rate functions (or, more concretely, by the associated selection dynamics). In this paper, we use the standard n -population replicator dynamics. Their definition as well as general definitions related to growth rate functions are recalled in the next subsection, along with an explanation of why it seems reasonable to use the replicator dynamics in our framework.

B.2 Growth rate functions and replicator dynamics

Mathematically, the change of the share of different strategies in the population over time can be captured by *growth rate functions*. A growth rate function g assigns to each population state x , player population i and pure strategy h available to a player from population i the growth rate $g_{ih}(x)$ of the associated population share x_{ih} . More concretely, we assume that $\dot{x}_{ih} = g_{ih}(x)x_{ih}, \forall i \in \{1, \dots, n\}, h \in \{1, \dots, S_i\}, x \in \Theta$, where \dot{x}_{ih} denotes the time derivative of x_{ih} and Θ is the polyhedron of mixed strategy profiles of all subpopulations (which equals the set of all possible population states). The following definitions are taken from Weibull (1997).

Definition 1. *A regular growth-rate function is a Lipschitz continuous function $g: X \rightarrow \mathbb{R}^{S_1 + \dots + S_n}$ with an open domain $X \subset \mathbb{R}^{S_1 + \dots + S_n}$ containing Θ , such that $g_i(x) \cdot x_i = 0$ for all population states $x \in \Theta$ and player populations $i \in \{1, \dots, n\}$.*

In this definition, $g_i(x)$ refers to the vector $(g_{i1}(x), \dots, g_{iS_i}(x))$ of growth rates of the strategies in subpopulation i . The condition $g_i(x) \cdot x_i = 0$ then makes sure that the overall size of the subpopulation does not change and only the shares of the different strategies are shifting. Before we go on, let us denote by e^i the i -th standard unit vector¹⁴. Moreover, we denote by $B_i(x) = \{h \in S_i: u_i(e^h, x_{-i}) > u_i(x)\}$ the set of pure strategies that earn above average payoff for player position i when the population state is x . In this notation, $u_i(e^h, x_{-i})$ describes the payoff of an individual belonging to subpopulation i from playing pure strategy h when the other subpopulations are in state x_{-i} . Likewise, $u_i(x)$ denotes the (expected) payoff of an individual from subpopulation i when the whole population is in state x .

Definition 2. *A regular growth-rate function g is weakly payoff-positive if for all $x \in \Theta$ and $i \in \{1, \dots, n\}: B_i(x) \neq \emptyset \Rightarrow g_{ih}(x) > 0$ for some $h \in B_i(x)$.*

This means that a weakly payoff-positive growth-rate function makes sure that for every subpopulation at least one strategy earning an above average payoff will

¹⁴This is the vector with a one at component i and with zeros everywhere else.

strictly grow. The most famous growth-rate function is probably the one associated to the standard n -population replicator dynamics.

Definition 3. *The standard n -population replicator dynamics are the selection dynamics given by $\dot{x}_{ih} = (u_i(e^h, x_{-i}) - u_i(x))x_{ih}$. They are associated to the growth-rate function $g_{ih}(x) = u_i(e^h, x_{-i}) - u_i(x)$.*

The replicator dynamics can be seen as the mathematical concretization of the evolutionary dogma *survival of the fittest*. The idea is simple: Once a strategy earns higher than average payoffs for a certain population, the share of the population playing this strategy will increase in the next periods while strategies with a lower than average payoff will become less popular or might even become extinct. The intuitive explanation for this can be evolutionary: Successful individuals will be able to feed more descendants and replicate while unsuccessful ones can not afford replication. In our application, less successful sports teams or trainers might look for a different profession or adapt their strategy so that the original strategy will disappear from the population in the long run. Hence, another explanation for the use of replicator dynamics is dynamic learning: Individuals observe what strategies work well and which do not and they are more likely to copy successful strategies than unsuccessful ones. It is straightforward to see that the replicator dynamics are weakly payoff-positive dynamics. More properties of the replicator dynamics can be found in Hofbauer and Sigmund (2003) and Sorin (2020).

B.3 Stable states in evolutionary games

In evolutionary game theory, we are interested in evolving systems of different populations, each of which may have different strategies at its disposal. A state in such a system can be described by a vector $x \in \mathbb{R}^{S_1-1+\dots+S_n-1}$, where n is the number of distinct subpopulations and S_i is the number of pure strategies from which the individuals of subpopulation i can choose.¹⁵ In Section 2.4.2, we study a setting with three (sub-)populations where each population has two pure strategies available. Consequently, a state in this system is given by a vector in \mathbb{R}^3 .

When we refer to $x(0)$ as the initial state of the system, we use $x(t)$ to denote the state of the system after time t . This state depends on the concrete (discrete or continuous) dynamics of the system. In the case of continuous dynamics, the

¹⁵The dimension of the vector space is due to the fact that the shares of the members of each subpopulation employing the different strategies must sum to 1.

evolution of the strategy shares of each subpopulation is described by the replicator dynamics through a differential equation. The Picard-Lindelöf theorem ensures that the trajectory $(x(t))_{t \in (0, \infty)}$ is indeed unique if the dynamic evolution rules are well-behaved. To simplify the notation, we denote by $\xi(t, x)$ the state of the system after time t , starting from state x . A state x^* is called a *steady state* if $\xi(t, x^*) = x^*$ for all $t \geq 0$. This means that once a system has reached a steady state, it will never leave it. Since a state corresponds to a point in \mathbb{R}^n , we will use the terms steady state and *stationary point* interchangeably. However, whether a system ends up in one steady state or the other depends on the stability properties of those states.

Definition 4. *A steady state $x \in \mathbb{R}^n$ is Lyapunov stable if every open neighborhood U of x contains an open neighborhood U^0 of x such that $\xi(t, x^0) \in U$ for all $x^0 \in U^0$ and $t \geq 0$.*

Definition 5. *A steady state $x \in \mathbb{R}^n$ is (locally) asymptotically stable if it is Lyapunov stable and there exists an open neighborhood U^* of x such that for all $x^0 \in U^*$, $\lim_{t \rightarrow \infty} \xi(t, x^0) = x$.*

Intuitively, Lyapunov stability requires that there be no push away from the state, while asymptotic stability requires a (local) pull toward the state. Asymptotically stable states are interesting because they can be a good predictor of what state a system will eventually settle on if it gets sufficiently close to that state at a given time.

Definition 6. *A steady state that is not asymptotically stable is called unstable. An unstable steady state that still has a lower-dimensional set of trajectories converging to it is called a saddle point. If all trajectories diverge from an unstable steady state, it is called a source or completely unstable.*

B.4 Relations between Nash equilibrium and evolutionary concepts

To conclude this brief overview over relevant elements of evolutionary game theory, we state some results linking the introduced concepts to Nash equilibria.

Definition 7. *A population state is called interior if and only if every strategy of every population has a positive share in this state.*

Theorem 4 (Weibull (1997), Theorem 5.2). *If the underlying growth-rate function g is weakly payoff-positive, then the following hold for a population state x :*

i) If x is an interior state and it is stationary, then x is a Nash equilibrium of the underlying game.

ii) If x is Lyapunov stable, then x is a Nash equilibrium of the underlying game.

Proposition 2 (Weibull (1997), Proposition 5.13). *A population state x is asymptotically stable in the standard replicator dynamics if and only if x is a strict Nash equilibrium¹⁶.*

C Data Collection

This section describes the collection process for the chess games used in Section 2.6. The first step was to search ChessBase’s online database for all games played by a certain player with the black pieces in 2018 and 2019. ChessBase is a company that produces chess software and maintains large databases of recorded chess games. Then, the games played in a shorter time control (rapid chess or blitz) were deleted from the data set. Finally, it was verified that the resulting set of games was complete by comparing the games with the FIDE-rated games on each player’s FIDE profile. FIDE is the international chess federation and provides ratings for all chess players based on their games played in official tournaments. A player’s FIDE profile shows the results of all the games a player has played in a given month.

Player	Number of games considered
Kirill Alekseenko	36
Fabiano Caruana	47
Anish Giri	27
Alexander Grischuk	24
Wang Hao	40
Ding Liren	44
Ian Nepomniachtchi	19
Maxime Vachier-Lagrave	29

Table 6: Number of games considered per player

¹⁶A *strict Nash equilibrium* is a Nash equilibrium in which the strategy used by each player constitutes the unique best response to the opposing players’ strategies.

The result of the comparison was that only two games played by Wang Hao were not recorded within this two-year period (which means the result is known, but not the moves played). Hence, these two games could not be included in the data set and the final number of games per player is given in Table 6.

Chapter 3

Congestion and Market Thickness in Decentralized Matching Markets*

JOINT WITH JULIAN TEICHGRÄBER

3.1 Introduction

Decentralized matching markets suffer from congestion. This has been known for a long time and many markets have been centralized to achieve more efficient market outcomes. However, there are also matching markets in which a centralization is not feasible for two possible reasons. First, it may be too difficult to elicit the complete preference ranking from all participants. Second, participants are heterogeneous and some may not have strong incentives to participate in a centralized matching market. A prime example of such a market is the labor market, where different firms operating in the same sector compete for the same (or at least a very similar) set of workers. What are the main sources of congestion in such a labor market? According to Roth (2018), congestion is “the accumulation of more time-consuming activities than can easily be accommodated in the time available.” In a labor market, these time-consuming activities include two main steps in the process. First, firms must conduct job interviews to assess the workers’ suitability for the firm. Since

*We are grateful to Christoph Schottmüller, Alexander Westkamp, Marek Pycia, Markus Möller, Lennart Struth, Max R.P. Grossmann and participants of the AMES 2022 in Tokyo as well as seminar participants in Cologne and Zurich for their valuable comments and suggestions. The remaining errors are our own.

time and resources are limited, firms must select who to interview among all applicants. Second, after receiving a job offer, workers must decide whether to accept or reject the offer, which is also time-consuming (especially if a worker receives multiple job offers from different firms).

In these congested markets, we study the impact of market thickness on market outcomes. In other words, we study how the problems caused by congestion change depending on market thickness. In particular, we examine whether there is inequality in how good and bad agents are affected when the market becomes thicker. We find that in a labor market, all firms and good workers lose when the market becomes thicker, while only bad workers benefit from the congested market structures in a thick market. In practice, market thickness can be designed by regulating the number of times per period that matches can be arranged. This was famously achieved for the economic academic job market, which now takes place only once a year, with many universities and junior researchers aspiring to be matched. In our model, we simulate a change in market thickness by multiplying the number of market participants on both sides by a natural number n .

As we are interested in two-sided matching markets in this paper, we will call the two sides of the market *firms* and *workers*, since this is a major application of our framework. Further applications include dating platforms¹ and the housing market² as well as college admissions. We model the congestion problem by assuming that a firm remains unmatched whenever the worker rejects the firm's job offer. In the real world, after a rejection by the firm's preferred job candidate (which can be very time-consuming), the firm could of course make an offer to another worker they have interviewed. However, since a lot of time may have passed, it is not unlikely that this worker will reject the offer because she already accepted another offer in the meantime. This happens even though she might have accepted an earlier offer from that firm, which is exactly the congestion problem. In our model, we limit the number of workers that each firm can observe and make an offer to. This can be seen as a second source of congestion according to the definition given in Roth (2018). The limitation is reasonable in real-world applications, since interviews are costly and firms' resources are limited. We assume that firms are

¹These dating platforms can be seen as centralized to begin with, but as people make offers independently of the platform, the same congestion problems can arise. However, the platform might control the suggestions that different people receive and in this way indirectly control the expected match value.

²Landlords usually have a preference for tenants with a secure and high income. In this sense, all potential candidates can be ranked according to their suitability as new tenants.

ranked by their quality and that each worker always accepts the offer from the highest firm that made him an offer. In this sense, workers are not strategic in our model. Firms, on the other hand, can be strategic in their offering decisions. We assume that only firms can make job offers and that they can only make an offer to a candidate they have interviewed. Both assumptions seem plausible in most labor markets. In our model, a very good firm might always choose to make an offer to the best worker they could find. We call this strategy the *myopic strategy* because it neglects the possibility that a better firm might have made an offer to the same worker. When there are few workers on the market, the probability that the best-screened worker was also seen by another firm is high. Consequently, the firm might decide to make a job offer to the second-best worker because the probability that the offer will be accepted is expected to be higher. Contrary to the myopic strategies, we compute the Bayesian Nash equilibrium strategies of the firms in which they take their ranking among all firms as well as the total number of workers into account. We compute the difference in expected match values for these two strategies and also study the effects of a changing market thickness on high or low quality firms or workers and equilibrium outcomes. Our result that all firms and good workers lose when the market becomes thicker, while only bad workers benefit from a thick market, holds whether firms use myopic or Bayesian Nash equilibrium strategies. The market participants who lose when the market becomes thicker are also the ones who would support centralization of the market if it leads to assortative matching of all participants.

The rest of this paper is organized as follows. Section 3.2 discusses the related literature. Section 3.3 introduces the model. Section 3.4 gives an example illustrating the firm's tradeoff in the hiring process. The market equilibria depending on firms' behavior are studied in Sections 3.5 and 3.6. The case in which firms adopt a myopic strategy is addressed in Section 3.5, while Section 3.6 examines the case in which firms adopt equilibrium strategies maximizing their respective expected match values. In Section 3.7, we discuss comparative statics to understand which parts of both market sides benefit more from the market becoming thicker. Section 3.8 concludes.

3.2 Related Literature

Our paper adds to the literature on congestion in matching markets (Roth and Xing (1997), Shimer and Smith (2001), Arnosti et al. (2014), Romanyuk (2017), Jagadeesan and Wei (2018), He and Magnac (2022)) by studying the impact of market thickness on congestion problems. Of the above papers, Arnosti et al. (2014) is similar to our work in terms of modeling sources of congestion: In their framework, an applicant observed by one firm is only available with a certain probability because he might have accepted another firm's offer during the application process. In our model, firms do not observe how many other offers an applicant has received either. This means that they have to estimate how likely it is that their desired candidate will accept their offer. Jagadeesan and Wei (2018) study the effect of signals in a congested job market. These signals allow workers to indicate a high acceptance probability following an offer from the firm receiving the signal. He and Magnac (2022) analyze application costs, which are comparable to firms' constraints on the total number of screens in our model. Both reduce the set of workers a firm can hire.

The role of market thickness in matching markets has been studied by Akbarpour et al. (2020), Baccara et al. (2020) and Loertscher et al. (2022) in the context of dynamic matching markets. In these (typically centralized) markets, the planner must decide when to make matches depending on market thickness and how easily the agents on the market can be matched. The market we study is not dynamic, meaning that the total number of agents on both sides of the market does not change over time. However, we can vary this number of agents to study the impact of a thicker market on market outcomes. Another key difference is that in our model we do not have a central planner, but consider a decentralized market with strategic participants at least on one side of the market.

In our model, we find that low-quality firms using the Bayesian Nash equilibrium strategy might make an offer to their worst-screened worker to avoid competing with better firms for a good worker. This is similar to the prediction of segmentation in matching markets found in Jacquet and Tan (2007). The intuitive reason for worse firms to make an offer to worse workers is that the expected acceptance probability of these workers is higher. A similar behavior was also reported in MacMillan and Anderson (2019) in the context of college admissions in the U.S. They find that U.S. colleges track potential applicants to determine who is most likely to accept an offer.

Che and Koh (2016) and Kadam (2015) are close to our paper, since they study the screening process in college admission and labor markets, finding that good applicants can suffer from the institutions' strategic behavior in the application processes. Chade et al. (2014) also examines the college admissions problem with a commonly known ranking of colleges. However, students are strategic in their setting, since they impose positive application costs on students. Another aspect in which our paper differs from Chade et al. (2014) as well as Che and Koh (2016) and Kadam (2015) is that we put a focus on the impact of a changing market thickness on the overall hiring process. Lee and Schwarz (2017) study the interviewing process on the labor market (without a focus on market thickness) and its consequences for unemployment. They find that unemployment is minimized when there is a perfect overlap in the set of workers interviewed by different firms in the sense that two firms either interview the same set of workers or no common workers at all. Coles et al. (2010) present the job market for new economists which was modified to achieve a higher market thickness. This also highlights a way in which a thick market can be implemented in practice: Limiting the times per year that the market is open and matches can be made.

3.3 The Model

3.3.1 Market participants and market thickness

We denote by $N \geq 2$ the total number of firms and M the total number of workers in the market. We assume that $M > N$ holds so that there are more workers than firms in the market. Each firm i has a quality level $q_i \in [0, 1]$ and each worker j has a skill level $s_j \in [0, 1]$. Firms prefer to be matched with workers with higher skill levels and workers prefer to be matched with firms with higher quality levels. More concretely, the utility level of matched market participants is given by the quality or skill level of their match. Unmatched market participants receive zero utility. Both the skill levels of the workers and the qualities of the firms are drawn i.i.d. from the uniform distribution on $[0, 1]$. The quality levels of firms are public knowledge. The skill levels of workers, however, are private information and each firm must make efforts to screen the skill level of particular workers. Each firm wants to fill one job position and each worker can accept at most one job offer.

We model market thickness by multiplying both numbers of market participants

in an existing market by a natural number $n \geq 1$. The parameter n measures the market thickness and higher values of n correspond to a thicker market. An intuitive justification for this notion of market thickness could be that n different copies of a certain market configuration are considered whose ex ante distributions of skills and qualities are the same, but the realizations may be different. We are interested in how this inclusion of more market participants (holding the ratio constant) affects market outcomes in congested markets. In reality, the emergence of thicker markets could be policy-driven: Think of a new highway or bridge connecting two cities or a legal restriction on the times per year that matches can be made. The reason we want to increase both sides of the market while maintaining their ratio is this: In these two-sided matching markets, it is rarely a disadvantage to have more participants on the other side of the market, while it is usually a disadvantage to have more participants on your side. Therefore, the effects of increasing only one side of the market are less interesting than increasing both sides simultaneously, since a positive and a negative effect interfere and the net effect is a priori unclear.

3.3.2 Screening and offering

Each firm that is on the market can screen the skills of two randomly selected workers and will subsequently make an offer to one of these two candidates. The limitation of two screens per firm can be seen as a reflection of the firms' limited resources in the hiring process. Both staff and time are scarce resources, which makes screening potential job candidates costly. One could argue that higher quality firms often have more resources and are able to screen more workers or receive better applications on average due to their good reputation. However, even if some firms are able to screen more workers, they still face the strategic uncertainty that we study in this paper: They do not know whether a specific worker has also been screened by a higher-quality firm and may therefore reject an offer made to her. Screening more skilled workers on average is *ceteris paribus* beneficial for a firm, but modeling this aspect would make the model much more complex. Moreover, it means that two different effects of higher firm quality on matching outcomes would have to be disentangled. For this reason, the above extensions are left for further research and are not addressed in this paper.

We assume that each worker can receive arbitrarily many offers. Each worker then looks at all the offers and chooses the firm with the highest quality. As mentioned earlier, all unmatched firms and workers receive zero utility. Note that

this assumption is how we model congestion in this paper. From a conceptual point of view, congestion in the labor market arises primarily because single workers hold multiple job offers at the same time and accepting or rejecting these offers is time-consuming. Consequently, firms have an incentive to make their job offers to workers who are likely to accept the offer, since a rejection leads to a lot of time being wasted in the hiring process. We neglect the possibility of a second round of offers, since this would not alter the strategic considerations of the firms and would only make the model less tractable.

3.3.3 Strategies of the firms

The matching game can have multiple outcomes depending on the players' behavior. Since we assume that the workers always accept the offer of the highest quality firm, only the firms can act strategically. We will examine the market outcomes for two different types of strategies chosen by each firm. First, we consider the myopic strategy, in which each firm always makes an offer to the screened worker with the highest skill level.

Second, we study the Bayesian Nash equilibrium (BNE) strategy, in which each firm makes the job offer decision based on its own quality and market thickness, as well as the equilibrium strategies of the better firms. If the firm has the highest possible quality, it will always make an offer to its best screened candidate, while if its own rank among the firms falls, it might prefer to make an offer to a lower ranked worker among the workers it screened. Since only the decisions (and screens) of higher quality firms affect a firm's assignment, a firm's Bayesian Nash equilibrium strategy need only be optimal given the strategies of the better firms. The explanation for this is as follows: The highest quality firm will always make an offer to its best screened worker because it can be sure that the offer will be accepted (by assumption, firms observe both their own quality and the quality of all the other firms in the market). The next best firm might also make an offer to its highest screen, and this worker will always accept the offer unless she has also received an offer from the best firm. In this sense, it is irrelevant to a firm what the lower quality firms' strategies are and its strategy only needs to be a best response to all the higher quality firms' strategies. We assume that any firm will always make an offer to the higher screen whenever it is indifferent between making an offer to its first or second screen. This makes the Bayesian Nash equilibrium unique.

3.4 An illustrative example

In this section, we illustrate the tradeoff that lower quality firms face in the hiring process. To do this, we assume that there are two firms and four workers on the market. It is common knowledge that firm 1 has a higher quality than firm 2. Consequently, every worker who receives an offer from both firms will always accept the offer from firm 1. The skill levels of the four workers are assumed to be 0.8, 0.6, 0.4 and 0.2. Note, however, that the firms do not know these specific values, but learn about two of these four skill levels only after the screening process. Ex ante, firms only know that each worker's skill is uniformly distributed and independent of the other workers' skill levels. Firm 1 will always make an offer to the screened worker with the higher skill level because it can be sure that this worker will accept the offer. The optimal behavior of firm 2 is more complex. Let $E(s_1|s_2)$ denote the expected match value of firm 2 after it makes an offer to a worker with skill level s_1 , based on the information that another worker with skill level s_2 is also available in the market. This second worker skill level corresponds to the second screen of firm 2. Let us now derive the optimal behavior of firm 2 in the two following cases:

- i) Firm 2 screened the workers with skills 0.8 and 0.6
- ii) Firm 2 screened the workers with skills 0.6 and 0.4

In case i), firm 2 compares the two values $E(0.8|0.6)$ and $E(0.6|0.8)$. For this, it has to consider the likelihood of firm 1 making an offer to either of the two screens. As there are six ways to screen two workers out of four, these two values can be computed as follows:

$$\begin{aligned} E(0.8|0.6) &= 0.8 \cdot \mathbb{P}(\text{Firm 1 makes no offer to worker 0.8}) \\ &= 0.8 \left(\mathbb{P}(\text{Firm 1 did not screen worker 0.8}) \right. \\ &\quad \left. + \mathbb{P}(\text{Firm 1 screened 0.8 and a better worker}) \right) \\ &= 0.8 \left(\frac{3}{6} + \frac{2}{6} \cdot 0.2 \right) = \frac{4}{5} \cdot \frac{3.4}{6} = \frac{13.6}{30} \end{aligned}$$

$$\begin{aligned}
E(0.6|0.8) &= 0.6 \cdot \mathbb{P}(\text{Firm 1 makes no offer to worker 0.6}) \\
&= 0.6 \left(\mathbb{P}(\text{Firm 1 did not screen worker 0.6}) \right. \\
&\quad + \mathbb{P}(\text{Firm 1 screened 0.8 and 0.6}) \\
&\quad \left. + \mathbb{P}(\text{Firm 1 screened 0.6 and a better worker, but not 0.8}) \right) \\
&= 0.6 \left(\frac{3}{6} + \frac{1}{6} + \frac{2}{6} \cdot 0.4 \right) = \frac{3}{5} \cdot \frac{4.8}{6} = \frac{14.4}{30}
\end{aligned}$$

From the above, it can be seen that firm 2 should make an offer to the lower screen 0.6 in this case. This is mainly due to the possibility that firm 1 might have observed the same set of workers, in which case it would make an offer to the more skilled worker and the worker with skill 0.6 would still be available to firm 2. Of course, this consideration becomes relevant for firm 2 only if the difference in skill levels between the two screens is relatively small.

In case ii), firm 2 will compare the following two values.

$$\begin{aligned}
E(0.6|0.4) &= 0.6 \left(\frac{3}{6} + \frac{2}{6} \cdot 0.4 \right) = \frac{3}{5} \cdot \frac{3.8}{6} = \frac{11.4}{30} \\
E(0.4|0.6) &= 0.4 \left(\frac{3}{6} + \frac{1}{6} + \frac{2}{6} \cdot 0.6 \right) = \frac{2}{5} \cdot \frac{5.2}{6} = \frac{10.4}{30}
\end{aligned}$$

This shows that in this case, firm 2 will make its offer to the higher screen, even though the possibility that firm 1 has screened the same set of workers still exists. The intuitive reason for this is that there are two cases in which firm 1 could have screened the worker with skill 0.6 and one of the two other workers whose skills firm 2 does not know. In these cases, the probability that the second screened worker has a higher skill is 0.4 (while in the case above it was only 0.2), which would result in firm 1 not making an offer to the worker with skill 0.6. This example illustrates that firm 2 (or any firm that is not the highest quality firm) must weigh the benefit of a higher acceptance probability for lower-skilled workers against the potential skill difference it foregoes by not *daring* to make an offer to the higher-skilled worker.

3.5 Myopic Strategy: Analysis

In this section, we characterize the market outcome when firms are using the myopic strategy. In particular, we compute the expected match values both for firms and for workers, depending on their quality or skill level.

3.5.1 Expected match value for firms

We want to compute the expected match value for a firm, depending on the number of firms N , the number of workers M and its own quality q . We denote this expected match value by $\mathbb{E}[v(N, M, q)]$.

Lemma 3. *The expected match value $\mathbb{E}[v(N, M, q)]$ of a firm with quality q on a market with N firms using the myopic strategy and M workers is given by*

$$\mathbb{E}[v(N, M, q)] = \sum_{f=0}^{N-1} (1-q)^f q^{N-1-f} \binom{N-1}{f} \int_0^1 2\tilde{s}^2 \left(\frac{\binom{M}{2} - 1 - (M-2)\tilde{s}}{\binom{M}{2}} \right)^f d\tilde{s}.$$

Proof. See Appendix A. □

3.5.2 Expected match value for workers

Lemma 4. *The expected match value $\mathbb{E}[v^w(N, M, s)]$ of a worker with skill level s on a market with N firms using the myopic strategy and M workers is*

$$\mathbb{E}[v^w(N, M, s)] = \sum_{k=0}^{M-1} s^k (1-s)^{M-1-k} \binom{M-1}{k} \sum_{i=1}^N \frac{i}{i+1} \binom{N}{i} \left(\frac{k}{\binom{M}{2}} \right)^i \left(1 - \frac{k}{\binom{M}{2}} \right)^{N-i}.$$

Proof. See Appendix A. □

3.6 Strategic Targeting: Bayesian Nash Equilibrium Analysis

In this section, we derive firms' BNE strategies and compute their expected match values in equilibrium.

3.6.1 Notation and derivation of the BNE strategy

From now on, we denote by f_i the firm that has the i -th highest quality of all firms. The workers each firm screened are identified by their skill level and we denote the *pair of screens* of firm f_i by $\mathcal{S}^i = (s_{i,1}, s_{i,2}) \in [0, 1]^2$. Without loss of generality, we assume that $s_{i,1} > s_{i,2}$ for all i (as the probability of two different workers having the same skill level is zero). The action set of firm f_i is given by $\mathcal{O}^i = \{o_1^i, o_2^i\}$, with o_j^i being the action of making an offer to screen $s_{i,j}$. A pure strategy of firm f_i is a mapping $\sigma_i : [0, 1]^2 \rightarrow \mathcal{O}^i$ and we denote by σ_{-i} the vector of pure strategies of all firms but firm f_i .

Definition 8. For each firm f_i and each of its screens $s_{i,j}$, we define the acceptance indicator $A_j^i \in \{0, 1\}$ of screen $s_{i,j}$ following an offer of firm f_i . This random variable³ takes the value 1 when screen $s_{i,j}$ would accept this offer from firm f_i and 0 when she would not accept it.

With the help of these random variables A_j^i , it will be possible to compute the expected match values of each firm. We will distinguish between the ex ante expected match value (the firm knows its quality and its ranking among all firms, but not the screened skill levels) and the interim expected match value (the firm knows both its quality and its screened skill levels). We denote the interim expected match value of firm f_i following action o_j^i and given its screens \mathcal{S}^i as well as other firms' strategies σ_{-i} by $V_j^i(\mathcal{S}^i) := s_{i,j} \mathbb{E}(A_j^i | \mathcal{S}^i, \sigma_{-i})$. In order to compute this expected value, we need the following definitions.

Definition 9. For a given firm f_i and one of its screens $s_{i,j}$, we denote by $\mathcal{F}_{i,j} = \{f_c | c < i, s_{i,j} \in \mathcal{S}^c\}$ the set of all higher quality firms who also screened $s_{i,j}$.

Definition 10. For a given firm f_i and one of its screens $s_{i,j}$, we associate an event e with a partition of $\mathcal{F}_{i,j} \cup \{f_i\}$ in the following way:

$$e = \{\mathcal{G}_{i,j}^1, \dots, \mathcal{G}_{i,j}^k\},$$

where the $\mathcal{G}_{i,j}^m$ are subsets of $\mathcal{F}_{i,j} \cup \{f_i\}$ defined by equivalence classes of the equivalence relation $f_r \sim f_l \Leftrightarrow \mathcal{S}^r = \mathcal{S}^l$. Without loss of generality, we assume that $f_i \in \mathcal{G}_{i,j}^1$. This means that an event e is unique up to relabeling of the sets $\mathcal{G}_{i,j}^2, \dots, \mathcal{G}_{i,j}^k$. The number k has to satisfy $k \leq \min\{i, M - 1\}$.

³The acceptance indicator is a random variable, since the screens of the other firms are random from firm f_i 's point of view.

Intuitively, the sets $\mathcal{G}_{i,j}^m$ are a partition of the firms up to f_i that have screened worker $s_{i,j}$ according to their respective second screens.⁴ Those firms that have the same second screen (meaning the screen different from $s_{i,j}$) are in the same set $\mathcal{G}_{i,j}^m \in e$. We refer to one such event e as a *competing firms' alternative options partition* (CFAOP). These CFAOPs will be needed in a central proof later on. The following example illustrates how these partitions can look like.

Example 2. *Let us consider the firm f_5 and one of its screens, e.g. $s_{5,1}$. Two (out of many more) possible events for this firm are given by $e_1 = \{\{f_1, f_3, f_5\}, \{f_2, f_4\}\}$ and $e_2 = \{\{f_5\}, \{f_2, f_3\}\}$. In the event e_1 , all firms f_1 through f_5 have screened the worker $s_{5,1}$, while firms f_1, f_3 and f_5 have even screened the same second worker. Firms f_2 and f_4 have also screened $s_{5,1}$ and the same second worker, but this worker is different from the second worker screened by firms f_1, f_3 and f_5 . In the event e_2 , the worker $s_{5,1}$ was screened by the firms f_2, f_3 and f_5 (and potentially lower quality firms, which is not relevant for firm f_5). The second screen of firms f_2 and f_3 is the same and different from the second screen of firm f_5 . Note that firm f_1 or f_4 could potentially also have screened $s_{5,2}$, the second screen of firm f_5 , but they are not in this partition set as they have not screened $s_{5,1}$ in the event e_2 .*

Definition 11. *For a set \mathcal{M} , let $\mathcal{E}(\mathcal{M})$ be the set of all partitions of \mathcal{M} .*

With the above definitions and notations, we are ready to derive the interim expected match value of a firm. As a first step, we compute the acceptance probability of a certain worker following an offer from firm f_i , conditional on its pair of screens \mathcal{S}^i and the vector of other firms' strategies σ_{-i} .

Lemma 5. $\mathbb{E}(A_j^i | \mathcal{S}^i, \sigma_{-i}) = \mathbb{P}(A_j^i = 1 | \mathcal{S}^i, \sigma_{-i}) =$

$$\begin{aligned} & \sum_{\mathcal{M} \in \mathcal{P}(\{f_1, \dots, f_{i-1}\})} \left[\left(\frac{2}{M} \right)^{|\mathcal{M}|} \left(\frac{M-2}{M} \right)^{i-1-|\mathcal{M}|} \sum_{e \in \mathcal{E}(\mathcal{M} \cup \{f_i\})} \left(\left(\frac{1}{M-1} \right)^{|\mathcal{G}_{i,j}^1|-1} \right. \right. \\ & \cdot \left. \left(\frac{M-2}{M-1} \right)^{|\mathcal{M}|+1-|\mathcal{G}_{i,j}^1|} \prod_{l=2}^{|\mathcal{M}|} \left(\left(\frac{1}{M-l} \right)^{|\mathcal{G}_{i,j}^l|-1} \left(\frac{M-l-1}{M-l} \right)^{|\mathcal{M}|+1-\sum_{h=1}^l |\mathcal{G}_{i,j}^h|} \right) \right. \\ & \cdot \left. \left. \prod_{f_h \in \mathcal{G}_{i,j}^1 \setminus \{f_i\}} \mathbb{1}_{\sigma_h(\mathcal{S}^i) \neq o_j^h} \prod_{l=2}^{|\mathcal{M}|} \left(\int_0^{s_{i,j}} \prod_{f_h \in \mathcal{G}_{i,j}^l} \mathbb{1}_{\sigma_h(s_{i,j}, s) = o_2^h} ds + \int_{s_{i,j}}^1 \prod_{f_h \in \mathcal{G}_{i,j}^l} \mathbb{1}_{\sigma_h(s, s_{i,j}) = o_1^h} ds \right) \right) \right] \end{aligned}$$

⁴Note that we are abusing notation slightly by identifying a worker with her skill level. Again, this should not lead to any confusion, since the probability that two workers have the same skill level is zero.

Proof. The idea of the proof is to consider all CFAOPs for firm f_i and its screen $s_{i,j}$ and compute the expectation conditional on each specific CFAOP. In the above notation, one CFAOP is an element of $\mathcal{E}(\mathcal{M} \cup f_i)$, for \mathcal{M} being a set of better firms than f_i (in other words, \mathcal{M} is an element of the power set $\mathcal{P}(\{f_1, \dots, f_{i-1}\})$) that have also screened worker $s_{i,j}$. The proof can be found in Appendix A. \square

Note that the formula derived in Lemma 5 depends only on the strategies of firms that are better than firm f_i . Hence, it is possible to iteratively derive each firm's BNE strategy in the following way: Firm f_1 will always use the myopic strategy as its BNE strategy, since this is clearly optimal for firm f_1 . For $i \geq 2$, Firm f_i will choose strategy σ_i defined by $\sigma_i(\mathcal{S}^i) = o_1^i$ if and only if

$$\underbrace{s_{i,1} \cdot \mathbb{P}(A_1^i = 1 | \mathcal{S}^i, \sigma_{-i})}_{V_1^i(\mathcal{S}^i)} \geq \underbrace{s_{i,2} \cdot \mathbb{P}(A_2^i = 1 | \mathcal{S}^i, \sigma_{-i})}_{V_2^i(\mathcal{S}^i)},$$

where the strategies σ_j of firms f_j with $j < i$ are given by the previously calculated BNE strategies of these firms (and the strategies of the worse firms in σ_{-i} have no effect on the calculation). This implies that firm f_i will choose $\sigma_i(\mathcal{S}^i) = o_2^i$ whenever the inequality does not hold. Here, we assume without loss of generality that firm f_i will make an offer to its first screen whenever the expected values from both offers are equal (which happens with probability zero). Hence, we can define the interim expected match value of firm f_i given its screens \mathcal{S}^i (and given that all firms use BNE strategies) by $\hat{V}^i(\mathcal{S}^i) = \max\{V_1^i(\mathcal{S}^i), V_2^i(\mathcal{S}^i)\}$.

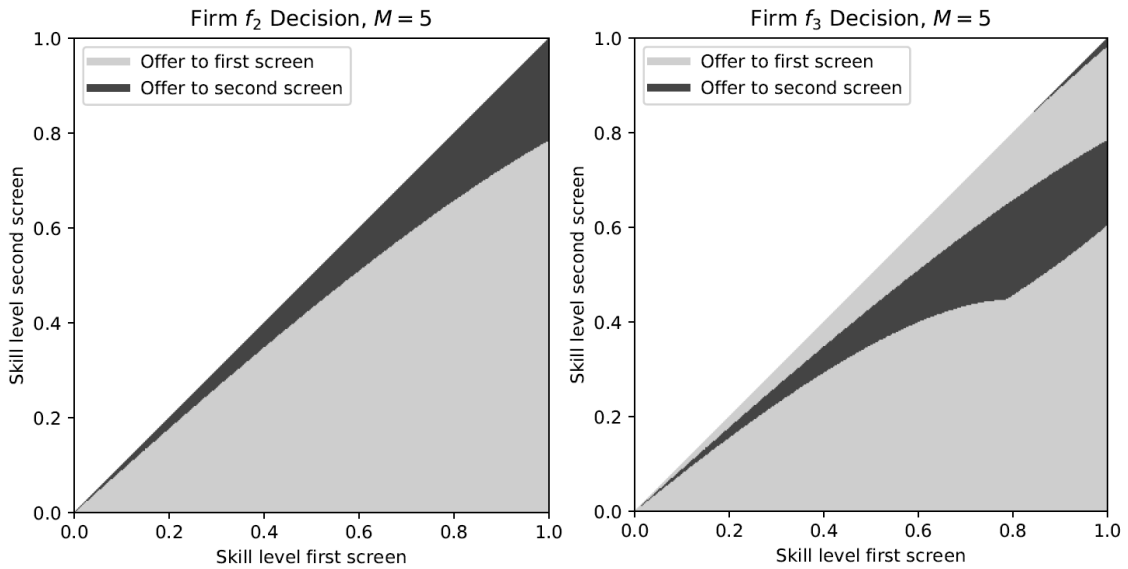


Figure 3.1: BNE strategies of firms f_2 and f_3 for $M = 5$ workers

Figure 3.1 shows the BNE strategies for firms f_2 and f_3 depending on their screens in a setting with 5 workers. In these plots, an interesting pattern can be observed. Firm f_2 will make an offer to its second screen only if both screened skill levels are close to each other. In these situations, firm f_2 is afraid of facing firm f_1 as a competitor for the highly skilled worker $s_{2,1}$. This also explains why the dark area becomes wider for higher values of $s_{2,1}$: In the cases where firm f_1 screened $s_{2,1}$, it becomes less likely that the second screen is better than $s_{2,1}$, which in turn makes it more likely that firm f_1 makes an offer to $s_{2,1}$. For firm f_3 , we find two areas in which the firm prefers to make an offer to its second screen. The upper area of those exists due to similar reasons as above. When both screens are very high and close to each other, firm f_3 is afraid to face either firm f_1 or firm f_2 as a competitor. This is why it makes an offer to the second screen, since the probability that this screen did not receive another offer is higher. Interestingly, for many other configurations close to the diagonal, firm f_3 makes an offer to its first screen. Intuitively, this is because firm f_3 asks itself how the better firms would act if they observed the same set of workers as firm f_3 . In this bright region close to the diagonal, firm f_1 would (as always) make an offer to its first screen, while firm f_2 makes an offer to the second screen. Consequently, both workers could be taken by a higher firm having screened the same two workers. This leads firm f_3 to making an offer to the higher screen, since this yields a higher match value in case the worker did not receive an offer from a better firm. Close to the upper border of the bright area of firm f_2 , firm f_3 also prefers to make an offer to the lower screen. This suggests that in this parameter region, firm f_3 is particularly afraid of competition from the part of firm f_2 who makes its offer to the first screen there. By making an offer to the second screen, firm f_3 can at least make sure that firm f_2 does not take its job candidate away in cases where they screened the same set of workers.

3.6.2 Ex ante expected match value of the firms

Starting from the interim expected match value of firm f_i , we can derive its ex ante expected match value by integrating the interim expected match value over all possible combinations of screens $s_{i,1}, s_{i,2}$. We denote the ex ante expected match value of firm f_i by \mathbb{V}^i and calculate it as follows:

$$\mathbb{V}^i = \int_0^1 2s_1 \int_0^{s_1} \hat{V}^i((s_1, s_2)) ds_2 ds_1,$$

as the distribution of the maximum of two independent and uniformly distributed random variables on $[0, 1]$ has density $2s$. In order to compare these expected values with firms' expected values calculated in Section 3.5, we express these values in terms of firms' qualities and not their overall ranking among all firms. This will also become useful when we compare them across different market thicknesses in Section 3.7. We denote these expected values by $\mathbb{V}(N, M, q)$ and compute them with the help of \mathbb{V}^i , for different values of i :

$$\mathbb{V}(N, M, q) = \sum_{i=1}^N (1-q)^{i-1} q^{N-i} \binom{N-1}{i-1} \mathbb{V}^i.$$

We will use this formula in Section 3.7 to study the impact of market thickness on market outcomes in congested markets.

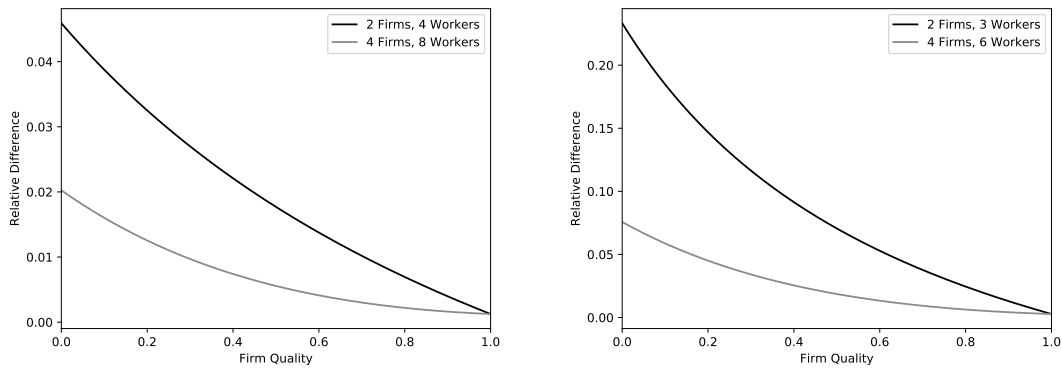


Figure 3.2: Relative difference in expected match values between BNE and myopic strategy

Figure 3.2 shows the relative difference in expected match values for a firm using the BNE strategy instead of the myopic strategy. More concretely, Figure 3.2 plots the values $\frac{\mathbb{V}(N, M, q) - \mathbb{E}[v(N, M, q)]}{\mathbb{E}[v(N, M, q)]}$ for different market configurations N, M and different firm qualities q . It shows that the benefit of using the BNE strategy is greater when the ratio of workers to firms is lower. This is intuitive, as it happens more often that a better firm takes away a good worker when the total number of workers is smaller. Another observation is that the relative advantage of the

BNE strategy over the myopic strategy is a lot bigger for firms with lower quality than for firms with a high quality. This is because high quality firms face less competition from better firms and making an offer to their best screen is less often a mistake. For low quality firms, it is more relevant to be strategic in their offering decisions, since the potential number of competing firms is higher.

3.6.3 Ex ante expected match values of the workers

We end this section by deriving the workers' ex ante expected match value $\mathbb{V}^w(N, M, s)$ depending on the number of firms N , number of workers M and their own skill level s . As a first step, note that the i -th highest quality firm has an expected quality of $\frac{N+1-i}{N+1}$ as all firm qualities are drawn independently from the uniform distribution on $[0, 1]$. A worker will be matched to the i -th highest quality firm if and only if this firm made an offer to her and all higher quality firms did not make an offer to her. As these events are independent from an ex ante perspective (before the firms' screens are determined), we only need to derive the ex ante probability that the i -th highest quality firm makes an offer to a worker with skill s . Let us denote this probability by $\mathbf{p}(i, s)$. Note that there are $\binom{M-1}{1} = M-1$ out of $\binom{M}{2} = \frac{M(M-1)}{2}$ cases in which the worker with skill s and one random additional worker are screened by firm f_i . The firm's BNE strategy σ^i determines for which skill values of the additional worker it will make an offer to the worker with skill s . As $\frac{M-1}{M(M-1)} = \frac{2}{M}$, the workers' ex ante expected match value when firms are using BNE strategies can be computed as follows:

$$\mathbb{V}^w(N, M, s) = \sum_{i=1}^N \frac{N+1-i}{N+1} \mathbf{p}(i, s) \prod_{j=1}^{i-1} (1 - \mathbf{p}(j, s)),$$

$$\text{with } \mathbf{p}(i, s) = \frac{2}{M} \left(\int_0^s \mathbb{1}_{\{\sigma^i(s,t)=o_1^i\}} dt + \int_s^1 \mathbb{1}_{\{\sigma^i(t,s)=o_2^i\}} dt \right).$$

3.7 Impact of Market Thickness: Comparative Statics

In this section, we study how a change in market thickness affects good or bad firms as well as good or bad workers in different ways. To do so, we assume that the number of firms changes from N to nN while the number of workers changes from M to nM , for a natural number $n > 1$. This leaves the ratio of firms to

workers unchanged while adding more participants to both sides of the market. We will compare how firms' ex ante match values following the myopic or the BNE strategy ($\mathbb{E}[v(N, M, q)]$ and $\mathbb{V}(N, M, q)$, respectively) react to these changes, depending on a firm's quality q . Likewise, we will conduct the same analyses for workers and their associated match values. We will use the expected match values and equilibrium strategies characterized in Sections 3.5 and 3.6. However, since these terms consist of large sums and the BNE strategy of a firm depends on all better firms' strategies, obtaining an analytical solution to our comparative statics questions is not feasible. This is why in this section, we will show illustrations for specific market configurations and provide analytical results that yield an intuition for why these findings are expected to hold more broadly. The Python code used to generate all the figures can be found under <https://github.com/mgramb/congestion-thickness>.

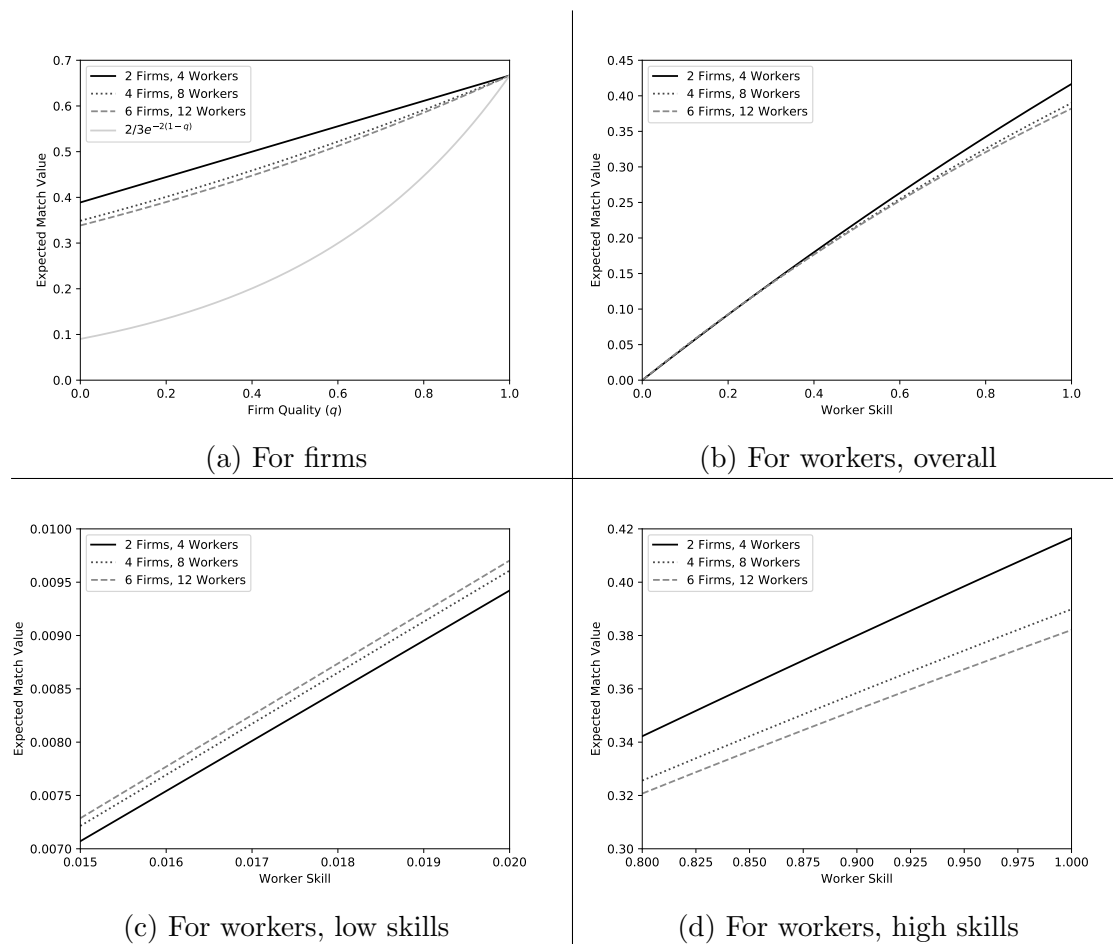


Figure 3.3: Expected match values when firms are using the myopic strategy

Figure 3.3 illustrates the implications of a higher market thickness when firms

are using the myopic strategy by showing three different market configurations. As can be seen in Figure 3.3a, an increase in n will always decrease the expected match value for a firm with given quality q . The reason for this is that the match value only depends on the (potential) number of better firms who make an offer to its first screened worker. When market thickness goes up, the probability of higher quality firms entering the market increases as well. At the same time, the probability that a single better firm has the same top-screened worker goes down due to the higher number of workers. However, as the match value is zero when there is only one better firm who makes an offer to the first screen, the negative effect of having more potential competitors outweighs the positive effect of an increased number of workers. The above effect is stronger for lower quality firms, since in expectation more entering firms will have a higher quality than them.

Figure 3.3b shows a different pattern for workers. Less skilled workers benefit from a thicker market (see Figure 3.3c). The reason is that they are more likely to be the top-screened worker of some firm and get an offer, since there are more firms and more potentially worse workers on the market. Highly skilled workers are losing when the market thickness goes up (see Figure 3.3d). The main reason for this is the following result.

Proposition 3. *A worker's probability of not being screened by any firm is increasing in market thickness n for all $n \geq 1$.*

Proof. See Appendix A. □

Proposition 3 explains why the expected match value of highly skilled workers goes down with an increase in market thickness when firms use myopic strategies: When these workers are screened by at least one firm, they often are the first-screened worker and get an offer due to their high skill level. When they are not screened by any firm, however, their match value is zero. Since the probability of this event increases in market thickness by Proposition 3, a higher market thickness is more disadvantageous for highly skilled workers.

When firms use BNE strategies, the impact of a thicker market is qualitatively the same as when myopic strategies are used. The effects on market participants' match values are illustrated in Figure 3.4.

The main reason why good workers suffer from a thicker market (see Figure 3.4d) is the same as when firms use the myopic strategy: The probability of not being screened by any firm increases with a thicker market. For low-skill workers

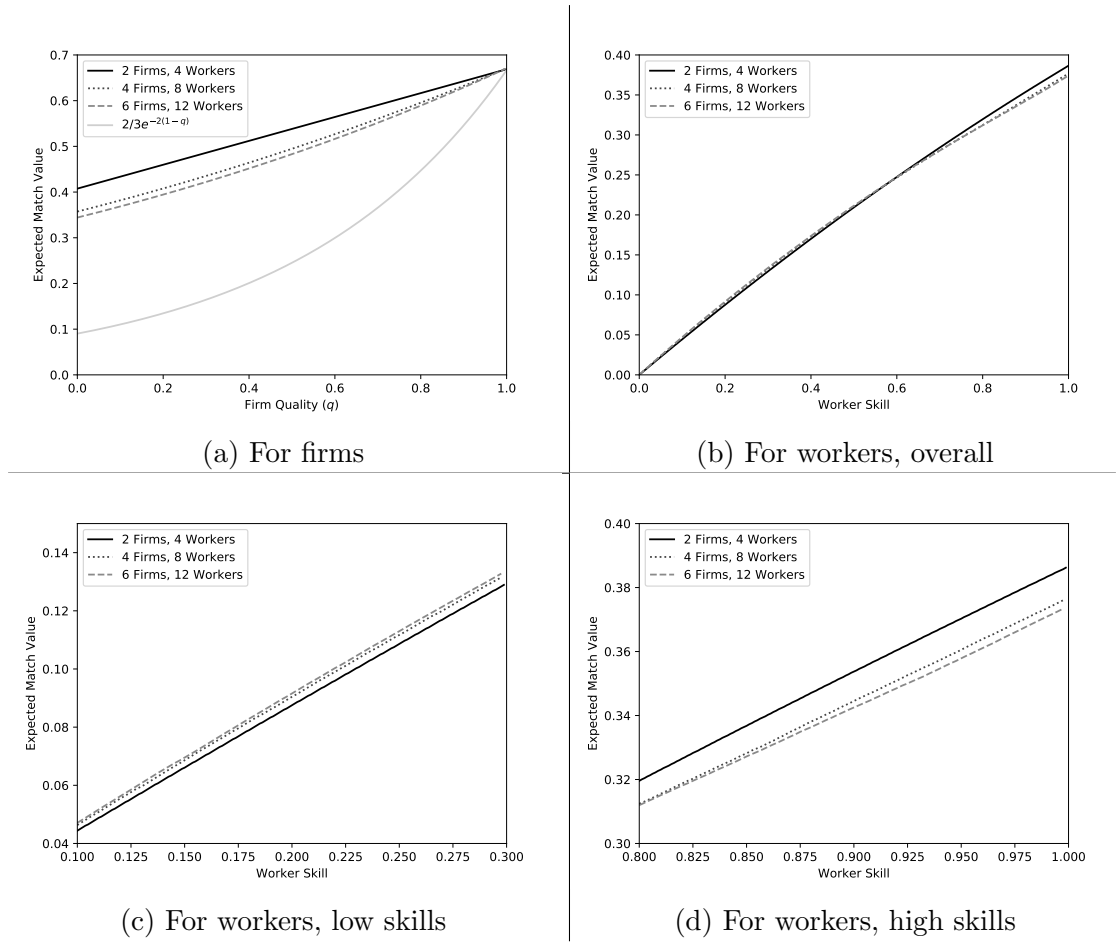


Figure 3.4: Expected match values when firms are using BNE strategies

(see Figure 3.4c), there is an additional reason for why they benefit from thicker markets when firms use BNE strategies. As firms take their own quality into account in their offering choice, some firms might make an offer to a low-skill worker simply because they are afraid of too much competition for their higher screen (caused by more better firms on average due to higher market thickness). This means that low-skill workers sometimes receive offers that they would not have received under myopic strategies.

We now present a limit result that justifies the trends in firms' match values in Figure 3.3a and Figure 3.4a.

Proposition 4. *For a given firm with quality q and a given market configuration with N firms and M workers, the probability of having no overlap in screened workers with any better firm (meaning that the intersections of the respective sets \mathcal{S}^i are empty) converges to the number $e^{-4(1-q)\frac{N}{M}}$ as market thickness n goes to infinity.*

Proof. See Appendix A. □

When there are many firms on the market, a single firm's match value depends to a large extent on the probability that it has no overlap in screened workers with higher firms. If this is the case, both screens would accept an offer from that firm. If, on the other hand, many better firms have at least one screen in common with the firm, the probability of either screen accepting a job offer is very low. Proposition 4 shows that this probability of having an empty overlap in screened workers with any better firm converges to a non-zero number for all firm qualities q and market configurations N, M . This number is increasing in q , which is natural as there are in expectation fewer better firms the higher the own quality q . The second implication of this convergence result is that firms' expected match values are bounded from below by the value $\frac{2}{3}e^{-4(1-q)\frac{N}{M}}$, as $\frac{2}{3}$ is the expected skill of the highest of two screens drawn from a uniform distribution on $[0, 1]$.⁵ This explains why the differences in expected match values for a given firm quality q become smaller the larger the market thickness n . The market configurations in Figures 3.3 and 3.4 are such that $M = 2N$, so the graph of $\frac{2}{3}e^{-2(1-q)}$ was added to Figures 3.3a and 3.4a.

Lastly, we study who would benefit from a centralized market and how this depends on market thickness. With a centralization of the market, we mean the creation of a central authority which is in charge of screening all the workers' values and making matches according to the preferences of all market participants. This would result in assortative matching in the sense that the best firm is matched to the best worker, the second-best firm is matched to the second-best worker, and so on.⁶

The main results of this analysis are illustrated in Figure 3.5. In this figure, it is assumed that firms use BNE strategies in the decentralized market. We see that all firms and the highly skilled workers would benefit from a centralized market while worse workers would lose. This is intuitive, since worse workers benefit from the random screening process in a congested decentralized market. Since screening is random, they are sometimes the best screened worker of a firm, even though their skill level is low compared to the average in the population. Moreover, firms are facing uncertainty about the number of better firms competing

⁵This implies that a firm using the myopic strategy guarantees itself an expected match value of at least $\frac{2}{3}e^{-4(1-q)\frac{N}{M}}$ in the limit $n \rightarrow \infty$.

⁶Such an assortative matching might be welfare-enhancing when there are complementarities in firms' qualities and workers' skills.

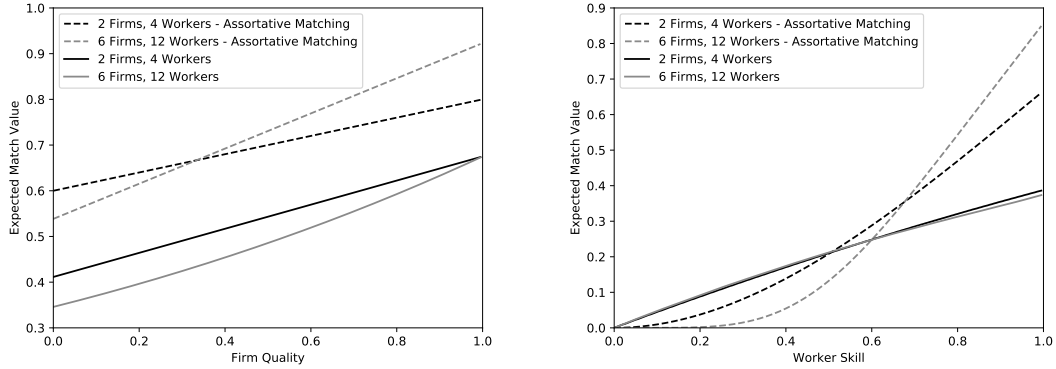


Figure 3.5: Assortative matching for firms and workers

for their screened workers on the decentralized market. Both of this implies that bad workers can sometimes be matched on a decentralized congested market, but never on a centralized market with assortative matching (where they do not belong to the best N among M workers).

The following lemma establishes a limit result that justifies the shape of workers' match values in Figure 3.5.

Lemma 6. *In the limit as $n \rightarrow \infty$, a worker's probability of being matched under assortative matching in a market with nN firms and nM workers converges to 1 if the worker skill s satisfies $s > 1 - \frac{N}{M}$ and it converges to 0 for $s < 1 - \frac{N}{M}$.*

Proof. See Appendix A. □

Lemma 6 implies that the workers' expected match values as a function of s in the limit as $n \rightarrow \infty$ can be described by a function that is zero for all $s < 1 - \frac{N}{M}$ and a linear function (connecting $(1 - \frac{N}{M}, 0)$ and $(1, \frac{N-1}{N})$) for $s \geq 1 - \frac{N}{M}$. A tendency toward a function of this type (with $\frac{N}{M} = \frac{1}{2}$) can already be seen in the two dashed lines representing worker values under assortative matching in Figure 3.5.

3.8 Conclusion

In this paper, we have studied a labor market with heterogeneous firms and workers. We find that it pays for (especially low-quality) firms to be strategic in the hiring process. Specifically, it is sometimes optimal for these firms to make an offer to a worker with a low skill level, even though they also screened another worker

with a higher skill level. We derived the Bayesian Nash equilibrium strategies for each firm and examined how the expected match values of firms and workers change in response to a change in market thickness. As the market becomes thicker, all firms and good workers lose as the problems created by congestion in the decentralized matching market become larger. Bad workers benefit from thicker markets because they are hired more frequently than in thinner markets. The results have implications on policy measures, as market thickness can be regulated, allowing policy makers to make more or less skilled people relatively better or worse off, depending on their policy goals.

We also found that in a decentralized labor market with more workers than vacancies, all firms and the good workers would support a centralization of the market with a resulting assortative matching. These market participants are typically more numerous and might have a stronger lobby (cf. Friedman and Friedman (1980, 216–17)) than the bad workers who are in favor of decentralized congested markets. This explains why many thick markets in various countries, cities and institutions have been centralized over the years.

Appendix

A Proofs

Proof of Lemma 3. Of course, the expected value depends on several factors:

- among the N firms, the number f of firms with a higher quality than q
- the number of workers M
- the skill level \tilde{s} of the highest screened worker

We will now first compute the expected value of a match assuming that these variables are given and subsequently take the expectation over each of the possible realizations of these variables.

$$\mathbb{E}[v(N, M, q)|\tilde{s}, f] = \tilde{s}(1 - \tilde{p}(M, \tilde{s}))^f,$$

where $\tilde{p}(M, \tilde{s})$ is the probability that a firm with higher quality (if existent) has the same top-screened worker. This probability naturally depends on M and \tilde{s} as well. Since $M \geq 3$ by assumption, we can compute $\tilde{p}(M, \tilde{s})$ as follows:

$$\tilde{p}(M, \tilde{s}) = \frac{1 + \binom{M-2}{1}\tilde{s}}{\binom{M}{2}} = \frac{1 + (M-2)\tilde{s}}{\binom{M}{2}}.$$

This leads to $\mathbb{E}[v(N, M, q)|\tilde{s}, f] = \tilde{s} \left(\frac{\binom{M}{2} - 1 - (M-2)\tilde{s}}{\binom{M}{2}} \right)^f$. We will now iteratively use these conditional expectations to compute new conditional expectations based on less variables.

$$\mathbb{E}[v(N, M, q)|f] = \int_0^1 \mathbb{E}[v(N, M, q)|\tilde{s}, f] \rho(\tilde{s}) d\tilde{s}, \quad (1)$$

with $\rho(\tilde{s}) = 2\tilde{s}$ being the density of the maximum of two screens,

$$\mathbb{E}[v(N, M, q)] = \sum_{f=0}^{N-1} (1-q)^f q^{N-1-f} \binom{N-1}{f} \mathbb{E}[v(N, M, q)|f]. \quad (2)$$

Finally, plugging in the value given for $\mathbb{E}[v(N, M, q)|f]$ in (1) into equation (2) yields the desired result. \square

Proof of Lemma 4. The expected match value of a worker with skill level s , given N and M , depends (directly or indirectly) on the following variables:

- the number i of firms that rank the worker first
- the number k of workers (among the $M - 1$ remaining ones) that have a worse skill level

In the following, we will first determine the expectation assuming that both of those values are known and then iteratively determine the expectation conditional on less variables.

$$\mathbb{E}[v^w(N, M, s)|i, k] = \frac{i}{i + 1},$$

as this is the expectation for the highest of i i.i.d. random variables that are uniformly distributed on $[0, 1]$. Next, we need the number $\mathcal{P}(w, k) = \mathbb{P}$ (a particular worker is the top-screened worker of a firm, given there are w other workers among which k workers are worse than him). This probability is zero when there are three or more workers in total (note that $M \geq 3$ by assumption) and the worker is the worst ($w \geq 2, k = 0$). Furthermore, when there is at least one worse worker ($w \geq 2, k \geq 1$), the probability is given by $\mathcal{P}(w, k) = \frac{k}{\binom{w+1}{2}}$ as the firm needs to have screened him as well as one worse worker for him to be top-screened. All in all, we obtain $\mathcal{P}(w, k) = \frac{k}{\binom{w+1}{2}}$ in both cases. Using this, we can iteratively take the expectation over i and k .

$$\begin{aligned} \mathbb{E}[v^w(N, M, s)|k] &= \sum_{i=1}^N \frac{i}{i + 1} \mathcal{P}(M - 1, k)^i (1 - \mathcal{P}(M - 1, k))^{N-i} \binom{N}{i} \\ &= \sum_{i=1}^N \frac{i}{i + 1} \left(\frac{k}{\binom{M}{2}} \right)^i \left(1 - \frac{k}{\binom{M}{2}} \right)^{N-i} \binom{N}{i}, \\ \mathbb{E}[v^w(N, M, s)] &= \sum_{k=0}^{M-1} s^k (1 - s)^{M-1-k} \binom{M-1}{k} \mathbb{E}[v^w(N, M, s)|k]. \end{aligned}$$

Putting all of the above formulas together yields the desired result. \square

Proof of Lemma 5. Note that $\mathbb{E}(A_j^i | \mathcal{S}^i, \sigma_{-i}) = \mathbb{P}(A_j^i = 1 | \mathcal{S}^i, \sigma_{-i})$. According to the law of total probability, this value is equal to

$$\sum_{\mathcal{M} \in \mathcal{P}(\{f_1, \dots, f_{i-1}\})} \mathbb{P}(\mathcal{F}_{i,j} = \mathcal{M}) \sum_{e \in \mathcal{E}(\mathcal{M} \cup \{f_i\})} \mathbb{P}(e | \mathcal{F}_{i,j} = \mathcal{M}) \cdot \mathbb{P}(A_j^i = 1 | \mathcal{S}^i, \sigma_{-i}, e).$$

In this formula, the sets \mathcal{M} are all the possible configurations of firms with a higher quality than firm f_i that could have screened $s_{i,j}$. We then calculate the probability of every possible event conditional on each specific configuration as well as the desired probability conditional on the event. The factors in the above formula are further computed below.

$$\begin{aligned} \mathbb{P}(\mathcal{F}_{i,j} = \mathcal{M}) &= \left(\frac{M-1}{\binom{M}{2}} \right)^{|\mathcal{M}|} \left(\frac{M-2}{M} \right)^{i-1-|\mathcal{M}|} = \left(\frac{2}{M} \right)^{|\mathcal{M}|} \left(\frac{M-2}{M} \right)^{i-1-|\mathcal{M}|} \\ \mathbb{P}(e | \mathcal{F}_{i,j} = \mathcal{M}) &= \mathbb{P}(\mathcal{G}_{i,j}^1 | \mathcal{F}_{i,j} = \mathcal{M}) \prod_{l=2}^{|e|} \mathbb{P}(\mathcal{G}_{i,j}^l | \mathcal{F}_{i,j} = \mathcal{M}, \mathcal{G}_{i,j}^1, \dots, \mathcal{G}_{i,j}^{l-1}) \\ \mathbb{P}(\mathcal{G}_{i,j}^1 | \mathcal{F}_{i,j} = \mathcal{M}) &= \left(\frac{1}{M-1} \right)^{|\mathcal{G}_{i,j}^1|-1} \left(\frac{M-2}{M-1} \right)^{|\mathcal{M}|+1-|\mathcal{G}_{i,j}^1|} \\ \mathbb{P}(\mathcal{G}_{i,j}^l | \mathcal{F}_{i,j} = \mathcal{M}, \mathcal{G}_{i,j}^1, \dots, \mathcal{G}_{i,j}^{l-1}) &= \left(\frac{1}{M-l} \right)^{|\mathcal{G}_{i,j}^l|-1} \left(\frac{M-l-1}{M-l} \right)^{|\mathcal{M}|+1-\sum_{h=1}^l |\mathcal{G}_{i,j}^h|} \\ \mathbb{P}(A_j^i = 1 | \mathcal{S}^i, \sigma_{-i}, e) &= \left(\prod_{f_h \in \mathcal{G}_{i,j}^1 \setminus \{f_i\}} \mathbb{1}_{\sigma_h(\mathcal{S}^i) \neq o_2^h} \right) \prod_{l=2}^{|e|} \left(\int_0^{s_{i,j}} \prod_{f_h \in \mathcal{G}_{i,j}^l} \mathbb{1}_{\sigma_h(s_{i,j}, s) = o_2^h} ds \right. \\ &\quad \left. + \int_{s_{i,j}}^1 \prod_{f_h \in \mathcal{G}_{i,j}^l} \mathbb{1}_{\sigma_h(s, s_{i,j}) = o_1^h} ds \right) \end{aligned}$$

Putting all the above formulas together yields the desired result.⁷ \square

Proof of Proposition 3. Let us denote by $P(N, M)$ the probability that a given worker is not screened by any firm on a market with N firms and M workers. We compute

$$P(N, M) = \left(\frac{\binom{M-1}{2}}{\binom{M}{2}} \right)^N = \left(\frac{M-2}{M} \right)^N = \left(1 - \frac{2}{M} \right)^N,$$

which directly implies that $P(nN, nM) = \left(1 - \frac{2}{nM} \right)^{nN}$. Even though n is supposed

⁷Here, we would have to be careful if we allowed $M = N$, because for M workers, only $M - 1$ partition elements are possible.

to be a natural number, let us take the partial derivative of this term with respect to n in order to see when it is increasing in n .

$$\begin{aligned} \frac{\partial P(nN, nM)}{\partial n} &> 0 \\ \Leftrightarrow 2 + \ln\left(1 - \frac{2}{nM}\right)(nM - 2) &> 0 \end{aligned}$$

As the last inequality is always satisfied for $nM \geq 3$ and we have $M \geq 3$ by assumption, we conclude that $P(nN, nM)$ is increasing in n for all $n \geq 1$. \square

Proof of Proposition 4. We first want to derive the probability to have an empty overlap with all better firms, for a given firm quality q , number of workers M and number of firms N . Let us denote this probability by $\mathbb{P}(N, M, q)$. The second step will be to compute $\lim_{n \rightarrow \infty} \mathbb{P}(Nn, Mn, q)$. To start with the first step, we compute the probability to have an empty overlap in screened workers with one specific firm. Since the screens of each firm are random, this probability only depends on the total number of workers M and we denote it by $\tilde{p}(M)$. The empty overlap happens exactly when the second firm has not screened any of the two workers screened by the first firm. Hence,

$$\tilde{p}(M) = \frac{\binom{M-2}{2}}{\binom{M}{2}} = \frac{(M-2)(M-3)}{M(M-1)}.$$

From this, we fix a firm quality q and compute $\mathbb{P}(N, M, q)$ as follows:

$$\begin{aligned} \mathbb{P}(N, M, q) &= \sum_{i=0}^{N-1} \mathbb{P}(i \text{ firms have quality higher than } q) \tilde{p}(M)^i \\ &= \sum_{i=0}^{N-1} (1-q)^i q^{N-1-i} \binom{N-1}{i} \tilde{p}(M)^i = (q + (1-q)\tilde{p}(M))^{N-1} \end{aligned}$$

Before we compute the limit in question, it makes sense to note that

$$(Mn - 2)(Mn - 3) = (Mn - 2)((Mn - 1) - 2) = Mn(Mn - 1) - 2(2Mn - 1) + 4. \quad (3)$$

$$\begin{aligned}
\lim_{n \rightarrow \infty} \mathbb{P}(nN, nM, q) &= \lim_{n \rightarrow \infty} (q + (1 - q)\tilde{p}(Mn))^{Nn-1} \\
&= \lim_{n \rightarrow \infty} \left(q + (1 - q) \frac{(Mn - 2)(Mn - 3)}{(Mn - 1)Mn} \right)^{Nn} \cdot \lim_{n \rightarrow \infty} \left(q + (1 - q) \frac{(Mn - 2)(Mn - 3)}{(Mn - 1)Mn} \right)^{-1} \\
&= \lim_{n \rightarrow \infty} \left(q + (1 - q) \left(1 - \frac{2(2Mn - 1)}{(Mn - 1)Mn} + \frac{4}{(Mn - 1)Mn} \right) \right)^{Nn} \quad (4) \\
&= \lim_{n \rightarrow \infty} \left(1 + (1 - q) \left(\frac{-2(2Mn - 1)}{(Mn - 1)Mn} + \frac{4}{(Mn - 1)Mn} \right) \right)^{Nn} \\
&= \left(\lim_{n \rightarrow \infty} \left(1 + \frac{(1 - q) \frac{-4}{M}}{n} \right)^n \right)^N \quad (5) \\
&= \left(e^{(1 - q) \frac{-4}{M}} \right)^N = e^{-4(1 - q) \frac{N}{M}}
\end{aligned}$$

In the above calculation, (4) is due to (3) and (5) follows because $\frac{-4Mn}{(Mn-1)Mn}$ behaves asymptotically like $\frac{-4}{Mn}$ (and $\frac{2+4}{(Mn-1)Mn}$ vanishes in the limit). This concludes the proof. \square

Proof of Lemma 6. Let us consider a worker with skill s on a market with nN firms and nM workers. Under assortative matching, this worker will be matched if and only if he is more skilled than the nN -th highest skilled worker among the $nM - 1$ remaining workers. As all the worker skills are uniformly distributed, his skill must be higher than the $n(M - N)$ -th order statistic of $nM - 1$ independent uniformly distributed random variables. This value follows a $Beta(n(M - N), nN)$ -distribution (see e.g. Gentle (2009, 63)). Its mean is $\frac{n(M - N)}{n(M - N) + nN} = 1 - \frac{nN}{nM} = 1 - \frac{N}{M}$ and the variance is $\frac{n(M - N)Nn}{(Mn)^2(Mn + 1)}$. This variance converges to zero as $n \rightarrow \infty$, so the original probability converges to 1 if $s > 1 - \frac{N}{M}$, while it converges to zero if $s < 1 - \frac{N}{M}$. This concludes the proof. \square

Chapter 4

Anonymous or personal?

A simple model of repeated personalized advice*

JOINT WITH CHRISTOPH SCHOTTMÜLLER

4.1 Introduction

In many situations, consumers ask better-informed experts to guide their choices. This happens even in situations where experts may have preferences over consumer choices that do not match the consumers' preferences, and it happens even in situations where it is difficult for the consumer to accurately articulate his exact preferences. For example, a consumer might ask his bank's employees for financial advice. The bank employee typically receives a bonus if the consumer purchases a particular investment product and often different investment products result in different bonuses for the adviser. There is no reason to believe that the product with the highest bonus is also the one best suited for the consumer. Similar situations occur in other retail sectors, such as consumer electronics or even cars.

Another example is internet search. A consumer enters a search term and relies on the search engine's response. Since some links are sponsored, there is an incentive for the search engine to emphasize the sponsored links more than links that better fit the consumer's needs but are not sponsored. A third example would be a minister (or manager) asking a civil servant (subordinate) to draft a particular

*We gratefully acknowledge the financial support of the Center for Social and Economic Behavior (C-SEB) of the University of Cologne.

legislative act or decree. Even if the civil servant has no policy preferences of his own, he might be aware that a similar draft has already been written under a previous government and that handing that old draft to the minister would save him a lot of time and effort. Again, this old draft is unlikely to do exactly what the minister wanted to accomplish. As a final example, consider a physician-patient relationship. The patient describes his symptoms and the physician prescribes a medication. Given the lobbying efforts of the pharmaceutical companies, it is quite possible that the physician has a preference for a certain drug company or pharmaceutical product.

What do these examples have in common? A consumer asks an expert to help him make a choice, although he cannot be sure what the expert's preferences are. In none of the examples is there a direct payment from the consumer to the expert, which means that the consumer has little ability to provide the expert with the right incentives. Furthermore, the consumer's communication of his preferences is complicated (due to the complicated nature of the issue and the consumer's ignorance that leads him to seek advice in the first place) and the expert's task is difficult. In other words, even if the expert tried to help the consumer as best he could, there would be some likelihood of misunderstanding and error. In a static one-shot game, we should not expect useful advice in any of these situations: By the one-shot nature, an expert would optimally recommend the alternative that earns him the (highest) bonus, since the consumer has no way to punish this behavior. Knowing this, the consumer would then not even ask for advice as the recommendation would not be consistent with his preferences. However, the above examples do not usually resemble a one-shot game. Consumers repeatedly consult the same financial adviser, use the same search engine, work with the same subordinates or visit the same physician. Repeated interaction – one could call it “relationship building” – has two interesting features: First, it is well known in game theory that cooperative behavior can be sustained in repeated interactions, even if this behavior cannot be sustained in a static one-shot game. Therefore, meaningful advice might be possible because of the repeated nature of the advice situation. Second, the adviser could learn to interpret the consumer's wishes. That is, the adviser's ability to give fitting recommendations is likely to improve over time. This is because both the adviser and the consumer can observe how previous recommendations have played out, such as whether the consumer was satisfied with the product purchased (or tried to return it), whether the consumer clicked

on the recommended link (and stayed on the website or subsequently purchased something there), whether the draft was pushed forward or discarded or whether the patient was cured. The success or failure of the recommendation can be used to learn how to interpret future requests from the consumer.

It should be noted that the learning we have in mind is relationship-specific. In particular, prior learning would be of little use to the consumer if he decides to switch experts. Although the consumer might also learn how to express his wishes to some extent, most of the learning seems to be on the expert's side. This paper therefore focuses on a setting where only the expert learns, and attempts to answer several questions. The most basic question is whether an equilibrium with meaningful advice is possible. The answer, unsurprisingly, is yes. The expert will give partially useful advice in equilibrium because the consumer threatens to end the relationship (and therefore the expert's opportunity to collect bonuses) if he receives bad advice for a number of periods. The key question is whether the consumer will benefit from the expert's learning. This is unclear because the consumer's outside option is not affected by the expert's learning, i.e. the expert could counteract his improved ability to give the right recommendation by recommending the product for which he receives a bonus more often. It is shown that – under certain conditions – the consumer in a certain class of simple equilibria nevertheless benefits. The reason for this is a *value effect*. The more the expert learns about the consumer, the more valuable the consumer is to the expert in the sense that the expected discounted bonus stream from that consumer is higher. The expert will lose more if the consumer ends the relationship and is therefore more inclined to give good advice to avoid exactly that. This leads to a testable prediction: The probability that a relationship will end now given that it has not already ended is lower the longer the relationship lasts.¹

The result that consumers benefit from expert learning provides a natural explanation for a puzzle that has emerged in the literature on privacy. People do not take even simple measures to anonymize their online activities. For example, most users use a search engine like Google directly, rather than using an anonymized service that redirects their search queries through another server before forwarding them to Google (and thus anonymizing them).² Privacy advocates emphasize that

¹More precisely, the probability that the relationship will end this period is lower than it was m periods ago, where $m \in \mathbb{N}$ is a number defined by the consumer's equilibrium strategy.

²There are many easy-to-use services of this type, such as <https://www.startpage.com> or <https://www.privatesearch.io>.

the more information the search engine has about a user, the greater the potential for exploitation (a simple exploitation method would be to display more sponsored links). The model shows that this is not the only effect. Due to the value effect, consumers also benefit from the search engine’s learning. Staying anonymous can lead to lower consumer surplus in the model of this paper. This also explains why consumers might prefer to get advice from the same person, such as having the same financial adviser at their bank whenever they go there, or staying with the same physician instead of switching every time they fall ill.

The rest of this paper is organized as follows: Section 4.2 discusses related literature. Section 4.3 presents the model and the analysis is performed in Section 4.4. Section 4.5 deals with welfare and anonymization. Most proofs of our theoretical results can be found in Appendix A. To compare our theoretical findings with real-world behavior, we conducted a laboratory experiment. The key results from this experiment are reported in Section 4.6. Section 4.7 discusses the results of this paper, Section 4.8 concludes.

4.2 Related literature

The consumer-expert relationship we study can be reinterpreted as a relationship between a principal and a noisily informed agent. In this sense, our work is naturally related to the cheap talk literature started by Crawford and Sobel (1982) and surveyed in Sobel (2013) and Blume et al. (2020). The fact that repeated interaction can be beneficial despite the lack of commitment is reminiscent of the literature on relational contracting started by Bengt Holmstrom (Holmström, 1978, 1982). There are two notable differences. First, most of the cheap talk literature is either static or deals with reputation concerns (Sobel, 1985; Benabou and Laroque, 1992; Park, 2005). Reputation issues are not addressed in the context of this paper but are addressed in Schottmüller (2019), where a similar model is used, but it does not allow for learning by the expert. Second, and more importantly, the cheap talk literature deals with a different misalignment of preferences. Typically, there is a one-dimensional decision and the expert is biased in one direction, e.g. he prefers slightly higher decisions than the decision maker. The structure here is different because the expert simply has a preferred option that is independent of the consumer’s optimal option. One implication of this structure is that no

meaningful advice is possible in a static setting, whereas this is obviously not the case in the cheap talk literature.

Li et al. (2017) analyze a repeated games setting in which the expert's and the principal's preferred projects are always distinct but the principal's project does not always exist. Moreover, there is always a default option that yields zero for all, and a disastrous project that yields $-\infty$ for all. Only the expert observes the identity of the projects, can communicate them and they are implemented if both expert and principal put effort into the same project. Our paper differs in two ways: First, both the expert's and the consumer's preferred option always exist and they can be equal. Second, the expert is not perfectly informed about the consumer's preferred option, but he receives a signal whose quality may increase over time.

The setting in Lipnowski and Ramos (2020) is probably closest to ours, since there the principal decides in each period whether to freeze the projects or delegate the project decision to the expert. The expert observes the quality of the project (high or low) and then decides whether to implement it or nothing, but the principal never learns the quality of the project. They study an intertemporal delegation rule to create incentives for the agent/expert and find that the agent represents the principal's interests only if dynamic incentives are provided. Our setting differs as (i) the agent has only noisy information and (ii) the principal does not "pause" the expert but fires him when he is dissatisfied with his advice. Furthermore, we focus on welfare dynamics in a class of simple equilibria.

Another related strand of literature is that on consumer protection in financial advice (Inderst and Ottaviani, 2012a,b, 2009). In these papers, the financial adviser is not only concerned with getting his bonus but also with the suitability of his advice. They focus on policy interventions that provide the adviser with the right incentives or payment schemes depending on whether consumers know the adviser is biased or not. In our framework, the expert is exclusively paid by his bonus and only cares indirectly about the suitability of his advice as the consumer threatens to leave him after receiving bad advice. Moreover, we model the improvement of the signal technology over time, while Inderst and Ottaviani mostly assume an exogenous and static signal.

An important application of our paper is search engines. Previous work on this market has focused mainly on ad pricing and auctioning (Edelman et al., 2007; Edelman and Schwarz, 2010; Eliaz and Spiegler, 2011) while we focus on the strategic interaction of search engine and user. More closely related is the literature

on privacy in the context of search engines. Computer science has provided ways to enable fully anonymous search through encryption even when the provider has no commitment power, see Byers et al. (2004) and Çetin et al. (2016). However, results on the benefits of personalization in internet search are ambiguous. On the one hand, already Spiekermann et al. (2001) argue that people value privacy protection but are not able to take the necessary means to meet this privacy protection goal. In the same vein, Acquisti et al. (2015) have demonstrated that people are unsure how to protect their data and what parts of their data are used for what purpose. They conclude that privacy protection should be regulated because naïve people will be harmed otherwise. We add to this literature by showing that even in the absence of naïveté it is unclear whether a user should allow personalization or not. In fact, users benefit from personalization in a certain class of simple equilibria. Experimental evidence shows that users value privacy to some extent (Tsai et al., 2011; Chellappa and Sin, 2005) and that sellers can benefit more than buyers from personalization (Hillenbrand and Hippel, 2019). On the other hand, some authors have shown that providing some personal data can benefit consumers, see Xu et al. (2007); Zimmer (2008).

4.3 Model

The model is a dynamic game with infinite time horizon. In each period, there are two options, one of which the *consumer* (C) must choose. One of the two options fits C 's needs and therefore gives him a payoff of 1 while the other option gives him a payoff of 0. C 's prior is that both options are equally likely to give him a payoff of 1.

The *expert* (E) receives a private and noisy signal about which option fits C 's needs. More precisely, E 's signal leads to a posterior in which one option has probability $p^k > 1/2$ to fit C 's needs and the other option has probability $1 - p^k < 1/2$ to fit C 's needs. Without loss of generality we call the option that is more likely to fit C 's needs option 1. The precision of E 's signal, p^k , is an element of a finite set $P = \{p^1, p^2, \dots, p^n\}$ with $1/2 < p^1 < p^2 < \dots < p^n < 1$. As E learns about C 's needs over time, precision improves in the following way: Whenever E recommends the option fitting C 's needs, precision improves from p^k to p^{k+1} (unless $p^k = p^n$ in which case precision remains unchanged).³

³The finiteness of P simplifies the exposition, but does not affect the results. As p^i cannot

The *expert's payoffs* are as follows: In every period, E has a *bonus option*. That is, E receives a bonus of 1 if he recommends this option to C while he receives a payoff of 0 otherwise. Each option has ex ante the same probability of being the bonus option and the identity of the bonus option is private information of E.⁴

The *timing* is as follows. In each period, E privately observes his signal and the identity of his bonus option. Then E recommends an option to C. C follows this recommendation and period payoffs realize. Both players observe whether the recommendation fits C's needs or not. Then, C decides whether to *end* or *continue* the game. If C ends the game, C receives an outside option V_O in the following period while E receives no payoffs in all future periods. If C continues, another period of the same game begins. Both players discount future payoffs with discount factor $\delta \in (0, 1)$. Needs and bonus option are assumed to be independent of each other and across periods.

In what follows, the word *hit* (*miss*) is used to denote the event that the recommendation fits (does not fit) the consumer's needs in a given period.

To make the problem interesting, C's outside option should be neither too attractive nor too unattractive. For example, V_O should be lower than the value the consumer would receive if he had a signal of precision p^n . If this was not satisfied, C would have the dominant strategy to end the relationship immediately. The outside option should also not be too low. More precisely, we assume that V_O is higher than the value C gets when E recommends his bonus option in each period. If this did not hold, there would be a unique perfect Bayesian equilibrium in which C always continues and E always recommends his bonus option. These two conditions are stated as

$$\frac{1/2}{1 - \delta} < V_O < \frac{p^n}{1 - \delta}. \quad (4.1)$$

Before turning to the players' strategies, let us discuss some modeling choices. We assume that the recommendation itself is payoff-relevant, i.e. E receives his bonus if he recommends the bonus option and C receives his payoff if the recommendation fits his needs. Put differently, there is no real decision by C whether or not to follow the recommendation. This is not unreasonable because C has uniform beliefs

increase above 1, learning must eventually flatten out in the sense that precision has to converge to an upper bound as i becomes large. Finiteness of P relieves us of the notationally burdensome task of taking limits in certain proofs and allows us to use backward induction right away.

⁴Note that more options for the expert would only make the analysis more tedious without really adding anything to the model, since the expert will only decide between his bonus option and the option he considers most likely to be the fitting option for the consumer.

and therefore cannot draw any inference from the recommendation itself about the likelihood that the recommendation fits his needs. Given that C has continued in the previous period and thereby asked for more advice, it seems logical to follow that advice. That is, there is no reason in the model to first ask for advice and then not follow it. It is also in line with certain applications, e.g. a consumer using a search engine will typically not refuse to click on a recommended link and most patients, as long as they can afford it, will take the prescribed medication. It is assumed that at the end of a period both C and E observe whether the given recommendation fitted the consumer's needs. In the examples mentioned earlier, this last assumption is reasonable: A salesperson will observe whether the consumer tries to return the product, the civil servant will observe whether his draft is pushed forward and the doctor will find out whether the patient recovers. In the search engine example, the search engine observes whether the link was clicked and – in the case of Google – to the extent that the target website uses GoogleAnalytics, `csi.gstatic`, GoogleAdSense or a GooglePlus button, Google also receives information about the user's subsequent behavior on the target website.

Note that the model assumes independence at several points. First, the bonus option is independent of the consumer's needs. This is one of the main differences to the cheap talk literature and appears naturally in the examples of the introduction. Second, there is some temporal independence in the sense that the consumer's needs and the bonus options are drawn independently in each period. One way to interpret this is that the requests of the consumer are unrelated, e.g. searching for an Italian restaurant in one period and for news in another period in the search engine example or suffering from different diseases in the patient-doctor example. In the financial advice example, the market environment and the set of available products may change from period to period.

As argued before, E gets to know the consumer better, so the precision of E 's signal should increase over time. Depending on the application, the precision might increase either after each interaction or after each hit or not at all. It seems realistic that a fitting recommendation tells more about a consumer's preferences than a non-fitting one. The assumption made here is that the precision increases with the number of past hits and that this increase is deterministic and commonly known by C and E . That is, no learning happens after misses. The special case of no learning at all will be analyzed later as a starting point.

It is worth noting that no meaningful advice would be possible if the game

was not infinitely repeated. Let us consider the static case. E has no incentive to recommend anything other than his bonus option. C therefore receives no information about which option is more likely to fit his needs. A similar situation emerges in a finitely repeated game. The static analysis applies to the last period. Since there is no meaningful communication in the last period, C should end the game after the penultimate period (regardless of history). Anticipating this, E will optimally recommend his bonus option in the penultimate period, regardless of what his signal is. Iterating this reasoning the game unravels and no meaningful advice is possible in any period. In the infinitely repeated game, the situation changes because future bonuses may motivate E to give truthful advice even if his bonus option is option 2. As there is no last period, there is no period in which these dynamic incentives break down.

What are the *strategies* of the players in this game? We assume that the players base their decision only on observed, payoff-relevant information. That is, C's decision depends only on the sequence of hits and misses in the previous periods.⁵ E has to decide in each period which option to recommend. His decision depends on his posterior belief, his bonus action and the history of hits and misses. In principle, his decision could also depend on the history of bonus options, but this possibility is neglected because his current and future payoffs do not depend on this information (neither directly nor indirectly as C's strategy cannot condition on this information, which C has not observed).

In the following, we employ two commonly used equilibrium notions and compare their outcomes. Both put further restrictions on strategies. First is the *Markov equilibrium*, where strategies condition only on the actions and information of the current period and a payoff-relevant state variable. The state variable is the current precision p^k . Consequently, E's strategy is a function $s_E: P \times \{1, 2\} \rightarrow [0, 1]$ that assigns a probability of recommending option 1 to every $p^k \in P$ and the identity of the bonus option. C's strategy is a function $s_C: P \times \{hit, miss\} \rightarrow [0, 1]$ that assigns a probability of continuing the game to every $p^k \in P$ and the success of this period's recommendation.

The second notion of equilibrium is (an extension of) *grim trigger*. C continues as long as the recommendations are hits. He ends the game if m consecutive recommendations are misses for some $m \in \mathbb{N}$. E plays a best response to this

⁵In principle, C observes the specific recommendations but since the option labels are not observed by him, he is unable to condition his strategy on these labels.

strategy. Of course, it remains to be shown that C's grim trigger strategy is a best response to E's best response, but this turns out to be straightforward unless V_O is too high.

4.4 Analysis

In the following, we study two classes of simple equilibria and demonstrate the welfare implications. In Markov equilibria, consumers do not benefit from learning. The logic is that the consumer's outside option does not improve when the expert learns and consequently the expert will not be willing to leave him a higher surplus. In a grim trigger equilibrium, we show that the consumer does benefit from the expert's learning. However, if we extend the grim trigger concept such that the consumer does not quit after the first bad advice but after, say, two consecutive bad advice, the consumer may even lose out (for some parameter values) due to the expert's learning.

4.4.1 Markov equilibrium

Note first that there is always a *babbling Markov equilibrium*. In this equilibrium, E will always recommend his bonus option and C will always stop the game. Clearly, these are mutually best responses given assumption (4.1). Therefore, the interesting question is not whether a Markov equilibrium exists but whether a Markov equilibrium with some information transmission exists. Before answering this question in general, it is useful to analyze the case without learning where the precision of E's signal remains constant. If C does not stop the game beforehand, this situation occurs after $n - 1$ hits in our model when E's signal has precision p^n .

4.4.1.1 Model without learning

Without learning, the state never changes and therefore a Markov strategy will only condition on this period's information/actions. That is, a strategy for E consists of two probabilities of recommending option 1 if (i) it is the bonus option and (ii) it is not. Similarly, a strategy of C consists of two probabilities of continuing: one in case of a hit and one in case of a miss.

In equilibrium, the probability of continuing is (weakly) higher in case of a hit than in case of a miss. Otherwise, E would have an incentive to give worst possible

advice, i.e. to always recommend option 2 if it is the bonus option (and possibly even if it is not) which, according to (4.1), automatically implies that C is better off ending the game.

Since the probability of continuing the game is higher in case of a hit than in case of a miss, it is optimal for E to recommend option 1 if option 1 is the bonus option. In this case the incentives of C and E are aligned. E's strategy can therefore be reduced to a probability α of recommending option 1 when option 2 is the bonus option.

While other equilibria can exist, we will focus on the case where C continues with probability 1 in case of a hit. Note that this provides the greatest incentive for E to be truthful. The restriction is not problematic: It is not hard to show that whenever a non-babbling Markov equilibrium exists, there exists a Markov equilibrium in which C continues with probability 1 in case of a hit. Furthermore, this is the equilibrium that Pareto dominates all other Markov equilibria. Under this constraint, C's strategy is simply a probability β of continuing in case the recommendation is a miss.

Denote E's equilibrium value, i.e. his discounted expected payoff stream at the start of a period (even before knowing the identity of the bonus option), by Π . If option 2 is the bonus option, E prefers recommending option 1 if

$$\begin{aligned} p\delta\Pi + (1-p)\beta\delta\Pi &\geq 1 + p\beta\delta\Pi + (1-p)\delta\Pi \\ \Leftrightarrow \beta &\leq \frac{(2p-1)\delta\Pi - 1}{(2p-1)\delta\Pi}. \end{aligned} \quad (4.2)$$

Denote C's equilibrium value by V and note that C is willing to continue only if $V \geq V_O$. Since this is independent of whether the current period's recommendation was a hit or a miss and since C continues for sure after a hit, C must either continue with probability 1 even after a miss, $\beta = 1$, or C must be indifferent, $V = V_O$. The former cannot happen in equilibrium: (4.2) cannot hold for $\beta = 1$ and E would therefore always recommend his bonus option. However, by (4.1), C would then

strictly prefer not to continue. Therefore, $V = V_O$ in equilibrium and consequently α has to be such that⁶

$$\begin{aligned} V_O &= \frac{1}{2}p + \frac{1}{2}(\alpha p + (1 - \alpha)(1 - p)) + \delta V_O \\ \Leftrightarrow \alpha &= \frac{2(1 - \delta)V_O - 1}{2p - 1}. \end{aligned} \quad (4.3)$$

By (4.1), $\alpha \in (0, 1)$. Hence, in an informative Markov equilibrium, E uses a mixed strategy and E is only willing to mix if (4.2) holds with equality. Given these equilibrium strategies one can determine the equilibrium values and obtain conditions for the existence of a non-babbling Markov equilibrium.

Proposition 5. *A non-babbling Markov equilibrium in the model without learning exists if and only if*

$$\frac{1 - \delta}{\delta} \leq \frac{4p - 3}{2}. \quad (4.4)$$

In such an equilibrium $V = V_O$ and $\Pi > 0$ and in the Pareto optimal Markov equilibrium α is given by (4.3) and $\beta = 1 - 1/[(2p - 1)\delta\Pi]$.

Note that condition (4.4) is more likely to be satisfied the higher p and δ are. Moreover, it implies $p \geq 0.75$, so the signal quality has to be quite high in order to guarantee the existence of a Markov equilibrium. Intuitively, this makes sense since the expert has to be incentivized to recommend option 1 in some cases even when it is not his bonus option. This will happen when the expert is more patient (high δ) or is reasonably sure to produce a hit (high p) in this case, such that the next period will be reached with higher probability.

4.4.1.2 Model with learning

Also in the model with learning, it is straightforward to see that E will always recommend option 1 when option 1 is the bonus option. As before, we will focus on non-babbling Markov equilibria in which C continues for sure in case of a hit. Strategies are therefore given by sets of probabilities $\{\alpha^k\}_{k \in \{1, \dots, n\}}$ and $\{\beta^k\}_{k \in \{1, \dots, n\}}$. The players' values, i.e. their expected discounted payoff streams at the start of a period with precision p^k , are denoted by Π^k and V^k . It follows from the previous subsection that such an equilibrium can only exist if (4.4) holds (for $p = p^n$). This

⁶As C is indifferent, we can determine his value $V = V_O$ by writing down the expected payoff stream if he continued for sure this period.

condition is necessary but not sufficient for the existence of a non-babbling Markov equilibrium and is therefore generalized below.

The first step is to show that in no period E will recommend option 1 regardless of the identity of the bonus option while C continues regardless of whether the recommendation is a hit or a miss. While this property is not surprising, it is also not straightforward: After all, recommending option 1 gives E a higher chance to move to the next highest precision and in principle it would be possible for this to motivate him to be truthful (if Π^{k+1} is sufficiently larger than Π^k).

Lemma 7. *In Markov equilibrium $\alpha^k = \beta^k = 1$ cannot hold for any k because E's best response to $\beta^k = 1$ is $\alpha^k = 0$.*

Lemma 8. *In every Markov equilibrium $V^k = V_O$ for all $k \in \{1, 2, \dots, n\}$.*

Lemma 8 implies E's strategy in Markov equilibrium. If the game reaches precision p^k with positive probability in a Markov equilibrium, then E has to mix such that C is indifferent between continuing and stopping. That is,

$$\begin{aligned} V_O &= \frac{1}{2}p^k + \frac{1}{2}(\alpha^k p^k + (1 - \alpha^k)(1 - p^k)) + \delta V_O \\ \Leftrightarrow \alpha^k &= \frac{2(1 - \delta)V_O - 1}{2p^k - 1}. \end{aligned} \quad (4.5)$$

Note that α^k , as given by (4.5), is in $(0, 1)$ by assumption (4.1). Consequently, E must be indifferent between recommending either option if the bonus option is option 2. This indifference condition determines β^k :

$$\begin{aligned} 1 + p^k \beta^k \delta \Pi^k + (1 - p^k) \delta \Pi^{k+1} &= 0 + p^k \delta \Pi^{k+1} + (1 - p^k) \beta^k \delta \Pi^k \\ \Leftrightarrow \beta^k &= \frac{(2p^k - 1) \delta \Pi^{k+1} - 1}{(2p^k - 1) \delta \Pi^k}. \end{aligned} \quad (4.6)$$

Note that Π^n is given by the stationary equilibrium value derived in the proof of Proposition 5. From this, Π^{n-1} and β^{n-1} can be obtained and by backward induction all other β^k and Π^k can also be obtained. A non-babbling Markov equilibrium exists if all such obtained β^k are in $[0, 1]$. The following proposition gives a necessary and sufficient condition for exactly this.

Proposition 6. *A non-babbling Markov equilibrium in the model with learning exists if and only if*

$$\frac{\delta^{n-2} 4p^n - 3}{1 - \delta 4p^n - 2} + \sum_{k=0}^{n-3} \delta^k \frac{4p^{k+2} - 3}{4p^{k+2} - 2} \geq \frac{1}{\delta(2p^1 - 1)}. \quad (4.7)$$

In this Markov equilibrium, $V^k = V_O$ and

$$\Pi^k = \frac{\delta^{n-k} 4p^n - 3}{1 - \delta 4p^n - 2} + \sum_{j=0}^{n-k-1} \delta^j \frac{4p^{j+k} - 3}{4p^{j+k} - 2} \quad (4.8)$$

for $k \in \{1, 2, \dots, n\}$, and α^k and β^k are given by (4.5) and (4.6), respectively.

4.4.2 Simple grim trigger strategies and m -equilibrium

Like most repeated games, the game described here has multiple perfect Bayesian Nash equilibria. We will now focus on a class of equilibria in which C employs the following particularly simple strategy: C continues the relationship unless the past $m \geq 1$ recommendations were misses. After m consecutive misses, C stops the game and consumes his outside option. Since this strategy is somewhat similar to the grim trigger strategies taught in introductory game theory, we will call this strategy a simple grim trigger strategy of length m or m -strategy for short. A perfect Bayesian Nash equilibrium in which C uses an m -strategy is called m -equilibrium.

When can an m -strategy be optimal for C? First, C must have a continuation value of at least V_O after any history that contains fewer than m consecutive misses. Second, continuing after m misses must result in a continuation value of at most V_O . The latter can be easily achieved: According to (4.1), it is optimal to end the game if E recommends his bonus option in all subsequent periods. In an m -equilibrium, continuing after m or more misses is clearly off the equilibrium path. Hence, the following off path beliefs of E will make this response optimal: If C has continued after m misses before, then E believes that C will end the game in the next period regardless of whether there is a miss or hit in the current period. Given this belief, it is clearly optimal to recommend the bonus option now. This implies that it is indeed optimal for C to end the game after m (or more) misses. These off path beliefs are not ruled out by perfect Bayesian Nash equilibrium or normal refinements.

Based on this off path construction, the following steps suffice to construct an

m -equilibrium. First, derive E's best response to C's m -strategy. Second, verify that C's continuation value on the equilibrium path is at least V_O . This implies that C's strategy is optimal as ending the game earlier always yields only V_O .

What is E's best response to an m -strategy? In a given period, E is always tempted to recommend the bonus option in order to secure a payoff of 1. The downside of this choice is that a miss is quite likely if the posterior belief that the bonus action fits C's needs is low. An additional miss brings E closer to the end of the relationship, stopping the bonus stream forever and therefore leading to a payoff of zero for E. It is immediate that E will always recommend option 1 if option 1 is the bonus option.

We denote the value of the expected discounted bonus stream after t consecutive misses, when the signal strength is p^k , by Π_t^k . After $t - 1$ consecutive misses, it is optimal for E to recommend option 1 instead of the bonus option (in case the two are not identical) if the following relation (4.9) holds.

$$\begin{aligned} p^k \delta \Pi_t^k + (1 - p^k) \delta \Pi_0^{k+1} + 1 &\leq p^k \delta \Pi_0^{k+1} + (1 - p^k) \delta \Pi_t^k & (4.9) \\ \Leftrightarrow \frac{1}{\delta(2p^k - 1)} &\leq \Pi_0^{k+1} - \Pi_t^k \end{aligned}$$

Note that in an m -equilibrium $\Pi_m^k = 0$. Consequently, E – for a given k – is most inclined to give good advice after $m - 1$ misses. The following lemma verifies a more general result: Π_t^k is decreasing in the number of misses t which implies that E becomes more eager to give good advice as the number of misses increases. Furthermore, E benefits from learning in the sense that Π_0^k is increasing in k .

Lemma 9. *In every m -equilibrium, Π_0^k is increasing in k and Π_t^k is decreasing in t .*

Lemma 9 has a direct implication for E's strategy in an m -equilibrium: As Π_t^k is decreasing in t , (4.9) is more likely to be satisfied for higher t (fixing k). Thus, for a given precision p^k , E will recommend the bonus option if t is low and option 1 if t is sufficiently high (in case the two do not coincide). This result is stated as a corollary for further reference.

Corollary 1. *In every m -equilibrium, E uses a precision dependent cutoff strategy. That is, E recommends the bonus option if the number of consecutive misses t with signal strength p^k is strictly below some threshold $l^k \in \{0, 1, \dots, m\}$ and recommends option 1 otherwise.*

Note that both the case $l^k = 0$, corresponding to E always recommending option 1, and the case $l^k = m$, corresponding to always recommending the bonus option, are allowed. For $t \geq l^k$, E's value can be written as $\Pi_t^k = 1/2 + p^k \delta \Pi_0^{k+1} + (1 - p^k) \delta \Pi_{t+1}^k$. Keeping in mind that $\Pi_m^k = 0$ in an m -equilibrium, backward induction gives for $t \in \{l^k, \dots, m - 1\}$

$$\Pi_t^k = \sum_{j=0}^{m-t-1} \delta^j \left(\frac{1}{2} (1 - p^k)^j + p^k (1 - p^k)^j \delta \Pi_0^{k+1} \right). \quad (4.10)$$

For $t < l^k$, E's value is $\Pi_t^k = 1 + \delta \Pi_0^{k+1}/2 + \delta \Pi_{t+1}^k/2$. Using the expression for $t \geq l^k$ above, iterating backwards yields for $t < l^k$

$$\Pi_t^k = \sum_{j=0}^{l^k-t-1} \delta^j \left(\left(\frac{1}{2} \right)^j + \left(\frac{1}{2} \right)^{j+1} \delta \Pi_0^{k+1} \right) + \sum_{j=0}^{m-l^k-1} \left(\frac{\delta}{2} \right)^{l^k-t} \delta^j \left(\frac{1}{2} (1 - p^k)^j + p^k (1 - p^k)^j \delta \Pi_0^{k+1} \right). \quad (4.11)$$

Using relation (4.10), we can derive the exact value of the threshold l^k :

Lemma 10. *The threshold l^k chosen by E in an m -equilibrium is given by*

$$l^k = \begin{cases} 0, & \text{if } \frac{1}{2} + p^k \delta \Pi_0^{k+1} \leq (1 - \delta(1 - p^k)) \left(\Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta} \right) \\ m, & \text{if } \Pi_0^{k+1} < \frac{1}{(2p^k - 1)\delta} \\ \max \left\{ 0, \left\lceil m - 1 - \frac{\ln \left(1 - (1 - \delta(1 - p^k)) \frac{\Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta}}{\frac{1}{2} + p^k \delta \Pi_0^{k+1}} \right)}{\ln(\delta(1 - p^k))} \right\rceil \right\}, & \text{else.} \end{cases} \quad (4.12)$$

Note that Lemma 10 also implies that $m > l^k$ always holds in the third case, since the logarithm in the numerator is negative (the negations of the first two conditions ensure that the term inside the logarithm is between 0 and 1). This justifies the following

Remark 1. *If $l^k = m$ for some k in an m -equilibrium, then also $l^i = m$ for all $i \in \{1, \dots, k - 1\}$. This follows directly from Lemma 10 as Π_0^{k+1} is increasing in k and $\frac{1}{(2p^k - 1)\delta}$ is decreasing in k .*

It is useful to first analyze the case without (further) learning which occurs after $n - 1$ hits.

4.4.2.1 Model without learning

For $k \geq n$, $\Pi_0^k = \Pi_0^n$ since there is no more additional learning. This implies that in an m -equilibrium, Π_0^n has to solve (4.11) with the same Π_0^n on both sides of the equation. Furthermore, l^n in this equation has to be optimal in the sense of (4.9). The following lemma implies that there exist unique Π_0^n and l^n satisfying these optimality conditions.⁷

Lemma 11. *E has a unique best response to C's m -strategy in the model without learning.*

Whether an m -equilibrium exists depends on C's outside option. If E's best response to C's m -strategy, as derived in the proof of Lemma 11, leaves C with a sufficiently high value after 0 misses, then an m -equilibrium exists.

Proposition 7. *An m -equilibrium in the model without learning does not exist if*

$$2p^n - 1 < \frac{1 - \delta + (\delta/2)^{m+1}}{\delta(1 - (\delta/2)^m)}. \quad (4.13)$$

If (4.13) does not hold, an m -equilibrium exists if and only if $V_O \leq \bar{V}_O$ for some \bar{V}_O satisfying (4.1).⁸

4.4.2.2 Model with learning

We start by deducing explicit formulas for the continuation value V_t^k of the consumer after t consecutive misses and with precision k . As $V_m^k = V_O$ in an m -equilibrium, the value for $t \in \{l^k, \dots, m-1\}$ can be derived by backward induction. We obtain

$$V_t^k = \sum_{j=0}^{m-t-1} (1 - p^k)^j \delta^j p^k (1 + \delta V_0^{k+1}) + (1 - p^k)^{m-t} \delta^{m-t} V_O. \quad (4.14)$$

⁷For uniqueness, we require the tie-breaking rule that E recommends option 1 if he is indifferent. Without this tie-breaking uniqueness is (only) generic.

⁸Note that for $m = 1$, the condition (4.13) reduces to $2p^n - 1 < \frac{1-\delta}{\delta} \Leftrightarrow p^n < \frac{3}{4} + \frac{1-\delta}{2\delta}$. This is exactly the existence condition (4.4) for a Markov equilibrium without learning.

Using this, we can also derive the value of V_t^k for $t < l^k$. It is given by

$$\begin{aligned}
V_t^k &= \sum_{j=0}^{l^k-t-1} \left(\frac{\delta}{2}\right)^j \frac{1}{2} (1 + \delta V_0^{k+1}) \\
&\quad + \left(\frac{\delta}{2}\right)^{l^k-t} \left(\sum_{j=0}^{m-l^k-1} (1-p^k)^j \delta^j p^k (1 + \delta V_0^{k+1}) + (1-p^k)^{m-l^k} \delta^{m-l^k} V_O \right).
\end{aligned} \tag{4.15}$$

Before we compute the expert's expected value Π_0 at the start of the game, we introduce some notation. In an advice relationship between a consumer and an expert, let $w = (w_1, \dots, w_{n-1})$ denote the vector of waiting times until the first, second, \dots , $(n-1)$ -th hit, where w_i denotes the number of periods in learning level i (with precision p^i) until the i -th hit occurred. In an m -equilibrium, $w_i > m$ implies that the consumer will fire the expert as he produced at least m consecutive misses. Hence, there are two types of possible histories in an advice relationship: First, those where the expert produced at least $n-1$ hits and reached the last precision level p^n . Second, those where the expert produced at least m consecutive misses before p^n was reached. We denote these two sets of histories by

$$\begin{aligned}
\mathcal{W}_n &= \{w = (w_1, \dots, w_{n-1}) \in \mathbb{N}^{n-1} | 1 \leq w_i \leq m \ \forall i \in \{1, \dots, n-1\}\} \text{ and} \\
\mathcal{W}_f &= \{w = (w_1, \dots, w_{j^*}) \in \mathbb{N}^{j^*} \text{ for some } 1 \leq j^* \leq n-1 | w_{j^*} = m+1, 1 \leq w_i \leq m \ \forall i < j^*\}.
\end{aligned}$$

The set of all feasible histories in an m -equilibrium is then given by $\mathcal{W} = \mathcal{W}_n \cup \mathcal{W}_f$. For any $w \in \mathcal{W}_f$, let us denote by $\text{len}(w)$ the dimension of the vector w . This value always corresponds to the learning level in which the expert gets fired because he produces m consecutive misses. We can now derive Π_0 .

Proposition 8. *In an m -equilibrium, let $(l^k)_{k=1, \dots, n}$ denote the vector of switching*

strategies for the expert, depending on the precision level. The expected value Π_0 of the expert at the beginning of the game is given by the formula

$$\begin{aligned} \Pi_0 &= \sum_{\bar{w} \in \mathcal{W}} \mathbb{P}(w = \bar{w}) \mathbb{E}(\Pi_0 | w = \bar{w}), \text{ with} \\ \mathbb{P}(w = \bar{w}) &= \begin{cases} \prod_{i=1}^{n-1} \left(\mathbb{1}_{\{\bar{w}_i \leq l^i\}} \left(\frac{1}{2}\right)^{\bar{w}_i} + \mathbb{1}_{\{\bar{w}_i > l^i\}} \left(\frac{1}{2}\right)^{l^i} (1-p^i)^{\bar{w}_i - l^i - 1} p^i \right), & \text{if } \bar{w} \in \mathcal{W}_n \\ \left(\frac{1}{2}\right)^{l^{\text{len}(\bar{w})}} (1-p^{l^{\text{len}(\bar{w})}})^{m-l^{\text{len}(\bar{w})}} * & \end{cases} \\ \mathbb{E}(\Pi_0 | w = \bar{w}) &= \begin{cases} \sum_{k=1}^{n-1} \left(\mathbb{1}_{\{\bar{w}_k \leq l^k\}} \sum_{h=0}^{\bar{w}_k - 1} \delta^h + \mathbb{1}_{\{\bar{w}_k > l^k\}} \left(\sum_{h=0}^{l^k - 1} \delta^h + \frac{1}{2} \sum_{h=l^k}^{\bar{w}_k - 1} \delta^h \right) \right) \delta^{\sum_{j=1}^{k-1} \bar{w}_j} \\ + \delta^{\sum_{j=1}^{n-1} \bar{w}_j} \frac{\sum_{g=0}^{l^n - 1} \left(\frac{\delta}{2}\right)^g + \sum_{g=l^n}^{m-1} \delta^g \left(\frac{1}{2}\right)^{g-l^n}}{1 - \sum_{g=0}^{l^n - 1} \left(\frac{\delta}{2}\right)^{g+1} - \sum_{g=l^n}^{m-1} \left(\frac{1}{2}\right)^{l^n} \delta^g (1-p^n)^{g-l^n} p^n}, & \text{if } \bar{w} \in \mathcal{W}_n \\ \sum_{k=1}^{l^{\text{len}(\bar{w})} - 1} \left(\mathbb{1}_{\{\bar{w}_k \leq l^k\}} \sum_{h=0}^{\bar{w}_k - 1} \delta^h + \mathbb{1}_{\{\bar{w}_k > l^k\}} \left(\sum_{h=0}^{l^k - 1} \delta^h + \frac{1}{2} \sum_{h=l^k}^{\bar{w}_k - 1} \delta^h \right) \right) \delta^{\sum_{j=1}^{k-1} \bar{w}_j} \\ + \delta^{\sum_{j=1}^{l^{\text{len}(\bar{w})} - 1} \bar{w}_j} \left(\sum_{h=0}^{l^{\text{len}(\bar{w})} - 1} \delta^h + \frac{1}{2} \sum_{h=l^{\text{len}(\bar{w})}}^{m-1} \delta^h \right), & \text{if } \bar{w} \in \mathcal{W}_f. \end{cases} \end{aligned}$$

The following result deals with the hazard rate, i.e. the probability that the expert is fired in a given learning level, conditional on having reached that level. More concretely, we denote by $HR(k)$ the probability of having m consecutive misses in an m -equilibrium after reaching precision level k . For the last precision level $k = n$, $HR(n)$ denotes the probability of being fired in this level without having scored a hit before (since the game is infinitely repeated, the probability of being fired in the last precision level is 1).

Proposition 9. *If $p^{k+1} \geq 1 - (1-p^k)^m 2^{m-1}$ holds for all k in $\{1, \dots, n-1\}$, then $HR(k)$ is decreasing in k in an m -equilibrium.*

Example 3. *To illustrate the above proposition, let us consider an $m = 2$ equilibrium with an initial precision of $p^1 = 0.51$. The subsequent precision levels that guarantee a decreasing hazard rate are given by $p^2 = 0.5198, p^3 \approx 0.5388, p^4 \approx 0.5746, p^5 \approx 0.6381, p^6 \approx 0.7380, p^7 \approx 0.8627, p^8 \approx 0.9623, p^9 \approx 0.9972$.*

Example 4. *Figure 4.1 shows the precision levels that ensure a decreasing hazard rate according to Proposition 9 for $m = 2, m = 3$ and $m = 4$.*

4.5 Welfare dynamics and anonymization

In this section, we discuss the dynamics of consumer surplus. Since the consumer's value equals his outside option regardless of the precision level in Markov equilib-

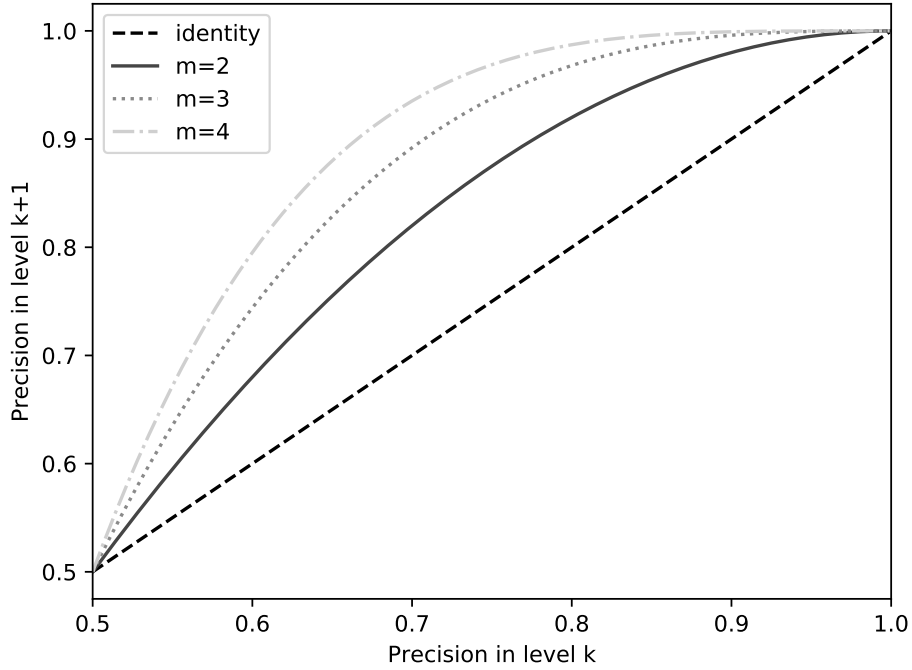


Figure 4.1: Sufficient learning jumps for a decreasing hazard rate

rium, the consumer does not benefit from learning in a Markov equilibrium. This is consistent with the argument that the expert can pocket all the benefit since the consumer's outside option (and therefore bargaining position) does not improve as the expert learns. However, analysis of m -equilibria shows that this logic may be flawed. Consider first the case of a classic grim trigger strategy, i.e. $m = 1$. The following proposition implies that consumers benefit from learning in this class of equilibria.

Proposition 10. *In an $m = 1$ equilibrium, V_0^k is strictly increasing in k and l^k is weakly decreasing in k .*

That is, the consumer can benefit for two reasons: simply because the expert's precision and therefore the advice quality improves but also because the expert's strategy can become more favorable over time. The intuition is that, by Lemma 9, the expert's profits are increasing in the precision level k (as long as the game continues). Therefore, as precision increases, he is more inclined to give good advice in order to reap the increasing future benefits.

To further illustrate the previous result and also to shed light on the dynamics

in m -equilibria for $m > 1$, we now consider the case of only two precision levels, p^1 and p^2 . We are primarily interested in the expert's choice of optimal thresholds l^1 and l^2 , since they determine the distribution of welfare between the consumer and the expert. First, we study the 1-equilibrium in which the consumer ends the relationship after the first miss. The following lemma shows that the expert's choice depends on how large δ is relative to p^1 and p^2 .

Lemma 12. *Let $n = 2$. If the consumer ends the game after one miss, the expert's advice choices l^1, l^2 are given by*

$$(l^1, l^2) = \begin{cases} (1, 1), & \text{for } \delta < \frac{1}{2p^2-1/2} \\ (1, 0), & \text{for } \frac{1}{2p^2-1/2} < \delta < \frac{1}{p^1+p^2-1/2} \\ (0, 0), & \text{for } \delta > \frac{1}{p^1+p^2-1/2} \end{cases} \quad (4.16)$$

The previous lemma implies that consumers benefit from learning in the $m = 1$ equilibrium: In equilibrium, l^2 (or more generally l^n for n rounds of learning) must equal 0. Otherwise, the consumer would be better off ending the advice relationship once the last precision level is reached, i.e. the $m = 1$ strategy would not be a best response. This implies that the cutoffs l^k are weakly decreasing in k in the $n = 2$ case and therefore advice improves in k for two reasons. First, the consumer can benefit from a lower l^k and thus a more honest advice strategy from the expert. Second, even if $l^1 = l^2 = 0$ and therefore the expert's advice strategy remains constant, the consumer benefits from learning as the signal technology improves.

While consumers benefit from learning in an $m = 1$ equilibrium, this is not necessarily the case in an $m > 1$ equilibrium. We illustrate this in the simplest possible case, i.e. only one round of learning ($n = 2$) and $m = 2$. In this case, we show that there are parameter values for which $l^1 = 0$ and $l^2 = 1$ in an $m = 2$ equilibrium. That is, the expert is less willing to give good advice after the signal technology improved. In our example, this change in expert strategy affects the consumer's payoff more than the improvement in signal technology and therefore the consumer's value will be lower at the beginning of a period with improved signal technology than at the beginning of the game.

Lemma 13. *Let $n = 2$. In an $m = 2$ equilibrium, $l^1 = 0$ and $l^2 = 1$ if and only if both*

$$\begin{aligned} & \frac{1 + \delta/4}{1 - \delta/2 - \delta^2 p^2/2} > \max \left\{ \frac{1 + \delta/2}{1 - \delta/2 - \delta^2/4}, \frac{(1 + \delta)/2 - \delta^2 p^2/2}{1 - p^2 \delta - (1 - p^2) \delta^2 p^2} \right\} \\ \text{and } & \frac{1}{2} + \frac{(1 - p^1) \delta}{2} + \frac{p^1 \delta (1 + (1 - p^1) \delta) (1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2} \\ \geq & \max \left\{ 1 + \delta/4 + \frac{(\delta/2 + \delta^2 p^1/2) (1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2}, 1 + \delta/2 + \frac{(\delta/2 + \delta^2/4) (1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2} \right\} \text{ hold.} \end{aligned}$$

Thus, an $m = 2$ equilibrium with $l^1 = 0$ and $l^2 = 1$ exists if the two inequalities above hold simultaneously. For $\delta = .98$, $p^1 = .85$ and $p^2 = .95$ both inequalities hold with strict inequality. Furthermore, for $V_O = 30$, we get $V_0^2 \approx 31.33 < 31.56 \approx V_0^1$, which proves the following result:

Proposition 11. *There exists an open set of parameters such that $V_0^1 > V_0^2$ in an $m = 2$ equilibrium.*

What is the intuition behind these results? Let us first consider Proposition 10. Knowing the consumer better means that the expert is better able to keep the consumer satisfied. The expert's value from giving good advice is higher when the signal is better because he is less likely to lose the consumer due to random errors. However, the value of recommending the bonus option does not depend on the signal technology. Thus, improvements in signal technology make it relatively more attractive to give good advice. Technically, the better the signal technology, the higher the continuation value of the relationship for the expert. This means that future payoffs and a continuation of the relationship gain in importance when the expert's signal technology improves and he is therefore more willing to give good advice. We call this the *value effect* of improved information and note that this effect is positive for the consumer.

Proposition 11 illustrates another dynamic effect that comes into play in more complicated equilibria. If the expert does not expect the consumer to end the advice relationship in case of a miss, it may be optimal for the expert to gamble: recommend the bonus option today and hope, in case of a miss, that recommending option 1 tomorrow will prevent the consumer from ending the relationship. The better the expert gets to know the consumer, the greater the incentive to gamble: The improved signal means that it is more likely that he will be able to provide a good recommendation if that is what is needed to keep the consumer tomorrow.

Put differently, the risk of ending the relationship is lower because the expert can be reasonably confident of providing a fitting recommendation “on the spot” if this is needed to keep the consumer. This *gambling effect* is negative for the consumer. In the example above, the gambling effect outweighs the value effect, so the consumer’s continuation value is higher when the signal technology is worse. Note that the gambling effect is not present in $m = 1$ equilibria, since in such an equilibrium the consumer ends the advice relationship immediately after the first miss.

We will now turn to the question of anonymization. The use of anonymized services makes relationship-specific learning impossible. For example, an internet search engine cannot personalize search results if the consumer uses an anonymized version of the search engine.⁹ In our model, anonymization corresponds to facing an expert who always remains at the precision level p^1 due to his inability to learn. Will the consumer benefit from anonymization? In a Markov equilibrium, the consumer surplus is always equal to the outside option, so anonymization has no effect. In an $m = 1$ equilibrium, on the other hand, anonymization harms the consumer: such an equilibrium exists only if E always recommends option 1 and in this case it is clear that the consumer would lose from anonymization.¹⁰ However, the consumer can benefit from anonymization in m -equilibria with $m > 1$. Consider the example above, on which Proposition 11 was based. An equilibrium with $m = 2$ also exists in the game where no learning is possible due to anonymization. In this anonymization equilibrium, $l = 0$ and the consumer value is $V_0^1 = 36.44$ which is larger than V_0^1 and V_0^2 in the equilibrium without anonymization. The intuition is that in this example learning leads to gambling, i.e. when precision equals p^2 , E is sufficiently confident that he can produce a hit on demand. Hence, he finds it optimal to recommend the bonus option in case the last recommendation was a hit. Without learning, the precision is too low to allow E to gamble and C benefits from sincere advice (albeit with a lower precision). This establishes the following result.

Proposition 12. *In Markov equilibrium anonymization neither harms nor benefits the consumer. In $m = 1$ equilibria the consumer always loses from anonymization*

⁹Anonymized versions of major internet search engines are widely available, see for example <https://www.startpage.com>.

¹⁰Note that the existence of an $m = 1$ equilibrium with anonymization implies that E will always recommend option 1 in the $m = 1$ equilibrium without anonymization: this follows directly from (4.9) and the facts that $\Pi_t^k = 0$ for $t > 0$ in $m = 1$ equilibrium and Π_0^{k+1} is weakly larger with learning than without.

while in $m > 1$ equilibria the consumer can benefit from anonymization for certain parameter values.

4.6 Experimental Design and Results

In this section, we present the design of our laboratory experiment and its main results. Additional results and robustness checks can be found in Appendix C.

4.6.1 Experimental Design

The experiment was conducted between December 2021 and February 2022 at the Cologne Laboratory for Economic Research, University of Cologne. We used the experimental software oTree (Chen et al. (2016)) and recruited participants via ORSEE (Greiner (2015)). The study was preregistered in the AEA RCT Registry (Gramb and Schottmüller (2022b)), its unique identifying number is: AEARCTR-0008682. Participants were randomly assigned to either the control group or the treatment group. In both groups, participants first read the instructions for their group, see Appendix B (in German), and answered a set of incentivized control questions. Then, players were randomly assigned the role of expert or consumer. Framing of roles was neutral in instructions and experiment. Subsequently, they played ten supergames (seven supergames in the pilot session in December, which was a treatment group session) of the game described in Section 4.3, each in their assigned role. After each supergame, each participant was randomly matched with another participant with the opposite role for the next supergame. The discounting of payoffs in the experiment was simulated by an exogenous stopping probability. After each round of a supergame, the game was exogenously ended with a probability of 10%, corresponding to a value of $\delta = 0.9$ in our model. In the control group, experts had a constant signal strength of 0.82. That is, the control group can be interpreted as a setting in which advice is given anonymously and therefore learning is not possible. In the treatment group, the first precision level was also 0.82 and precision was increased by 0.02 after each hit up to a maximum precision of 0.9. That is, the treatment group represents a setting in which personalized advice and incremental learning is possible. Once a consumer decided to end the game in either group, he immediately received a payoff of 5 points (while the expert's bonus and the consumer's payoff in case of a hit were both 1 point). It should be noted that in our model the outside option is paid out

at the beginning of the next period (since it always exists). Thus, the payout of 5 points after firing the expert corresponds to an outside option of $V_O = \frac{5}{\delta} = \frac{50}{9} = 5.\bar{5}$. After all supergames were completed, one supergame was randomly selected for each participant and the points earned there were paid out (with one point being worth 1€). At the end, participants were asked incentivized questions eliciting their risk attitude and completed a non-incentivized survey about trust attitude, age, gender and faculty. Additionally, each participant was paid a show-up fee of 4€. Participants' total payments ranged from 4€ to 25€. One session lasted between 29 and 56 minutes. There were seven sessions with a total of 156 participants in the treatment group and four sessions with a total of 98 participants in the control group. No participant attended more than one session.

4.6.2 Results

The main outcomes we are interested in are advice quality and consumer welfare in both groups. Let us start with advice quality. We measure this as the share of good advice given by the expert (in terms of the recommendation of option 1) in all situations where he faced a tradeoff (bonus option was option 2). Figure 4.2




Figure 4.2: Advice Quality in Control and Treatment group

shows that the advice quality in the treatment group is significantly better than in the control group. Hence, the potential increase in learning level incentivizes the experts to give better advice to retain consumers. As can be seen in Figure 4.3, this expert behavior leads to higher average consumer welfare in the treatment group, although the difference is not significant. A possible reason for the (only) small increase in consumer welfare in the treatment group is that consumers tend to distrust the expert more at higher learning levels. This can be seen in the firing rates in the treatment group.




Figure 4.3: Consumer Welfare in Control and Treatment group

In Figure 4.4, we see that the hazard rate¹¹ increases overall with learning level. Specifically, the hazard rate for precision levels p_2, p_3 and p_4 is significantly higher than for lower levels p_0 and p_1 . This could also drive the effect seen in Figure 4.2, where experts try to *convince* consumers not to fire them by giving them even more good advice. Note that this behavior does not contradict our theoretical predictions: The sufficient condition from Proposition 9 that the hazard rate decreases would require a level $p_1 \geq 0.9352$ for the value $p_0 = 0.82$. In the experiment, we set

¹¹In this general case, the hazard rate is simply the relative frequency with which consumers fire experts at a given learning level.

$p_1 = 0.84$, which is too small for the model to predict a decreasing hazard rate. One possible reason for this increase in hazard rate could be attribution of failure:




Figure 4.4: Hazard Rate per Learning Level in Control and Treatment group

At low learning levels, the consumer might attribute a miss to the expert's low signal strength. At high learning levels, it becomes increasingly likely that a miss is due to the expert's strategy to collect his bonus instead of giving good advice. This is then punished by the consumer who fires such experts. Interestingly, the hazard rate in the control group is significantly higher than the hazard rate for the first two learning levels in the treatment group. This suggests that consumers assume that the learning incentive has a positive effect on the relationship and do not fire the expert to establish such a long-term relationship.

4.7 Discussion

The results in Section 4.4 and 4.5 have implications for anonymization. Activists and experts alike recommend measures to preserve anonymity online. Although many of these recommendations are easy to follow, such as using an anonymized version of Google instead of Google itself, hardly any internet user follows them.

The above analysis indicates that consumers might be right not to anonymize: Personalized recommendations are more valuable not only from a total surplus perspective, but also from a consumer perspective in m -equilibrium if $m = 1$ (and often also if $m > 1$). The reason for this is simple. The more past usage data is available, the more valuable the customer is. The expert, e.g. Google, does not want to risk losing valuable customers. Hence, a customer enjoys better service when the expert can use past usage data from him. This theoretical finding is also supported by our experimental results: As we have shown in Section 4.6, experts give better advice in the treatment group where it is possible to learn from past interactions.

The same principle applies to other applications than Google and explains why long-term advisers are more valuable than short-term advisers. The m -equilibrium provides an interesting prediction for the hazard rate, i.e. the probability that a consumer will end the relationship after a certain number of hits if he has not already ended it. In an m -equilibrium, the hazard rate decreases over time when the change in signal quality between two learning levels is sufficiently high.

Of course, these results are subject to some caveats. The first is that the outside option of the consumer was held constant. If the outside option is an alternative expert, this could change. To give an example, say there are two experts and everyone agrees that Expert 1 is slightly more knowledgeable than Expert 2. The outside option then corresponds to getting advice from Expert 2. If everyone uses Expert 1, however, Expert 2 might be out of business and take up a different job. In the long term, the outside option might therefore decline and eventually drop below $1/(2 - 2\delta)$. In this case, the unique equilibrium is that the expert recommends his bonus action in each period and consumers would suffer. However, an m -equilibrium is not sensitive to lower outside options as long as the outside option remains above $1/(2 - 2\delta)$.

Another caveat, particularly in the context of anonymizing online activity, is that the model does not address potential extortion arising from abuse of data outside the advice relationship. According to the model, a customer benefits from personalized advice and a prerequisite for such personalized advice is that data about past interactions be stored. If this data gets into the hands of a third party, it could be used by that third party against the consumer; think health or financial records. Such third-party extortion is beyond the scope of this paper.

Another interesting result was given in Proposition 11 as it showed that too

much past data can also reduce the consumer's utility (although it is always higher than his outside option). Consequently, whether anonymization is optimal or not is ambiguous and depends on the particular equilibrium played as well as the model parameters.

4.8 Conclusion

In this paper, we have studied an expert-consumer relationship in which the expert gets to know the consumer over time and in this way can give better advice as the relationship progresses. We have shown that this learning opportunity can be beneficial to both the consumer and the expert by introducing m -equilibria as a generalization of simple grim-trigger strategies. Empirical evidence from our laboratory experiment suggests that experts do indeed give better advice when learning is possible. However, the consumer must be aware that too much learning on the part of the expert can be detrimental to consumer welfare. The choice of how much and what data to disclose is therefore a difficult one.

Appendix

A Proofs

Proof of Proposition 5. With $\beta = 1 - 1/[(2p - 1)\delta\Pi]$, it is straightforward to determine Π :¹²

$$\begin{aligned}\Pi &= \frac{1}{2} + (p + \beta(1 - p))\delta\Pi = \frac{1}{2} + \left(p + (1 - p) - \frac{1 - p}{(2p - 1)\delta\Pi}\right)\delta\Pi \\ \Leftrightarrow \Pi &= \frac{4p - 3}{(1 - \delta)(4p - 2)}.\end{aligned}$$

Plugging this back into (4.2) (with equality) yields

$$\beta^* = 1 - \frac{2(1 - \delta)}{\delta(4p - 3)}.$$

A non-babbling Markov equilibrium exists if $\beta^* \in [0, 1]$ which is the case if and only if

$$\frac{1 - \delta}{\delta} \leq \frac{4p - 3}{2}.$$

□

Proof of Lemma 7. Suppose $\beta^k = 1$ and distinguish between the two cases of either option 1 or option 2 being E's bonus option (both happen with probability $\frac{1}{2}$). E's value as a function of α is then

$$\begin{aligned}\Pi^k &= \frac{1}{2}(p^k\delta\Pi^{k+1} + (1 - p^k)\delta\Pi^k)(1 + \alpha) + \frac{1}{2} + \frac{1}{2}(1 - \alpha)(p^k\delta\Pi^k + (1 - p^k)\delta\Pi^{k+1} + 1) \\ &= \frac{1}{2}(2 - \alpha) + \frac{1}{2}\delta\Pi^{k+1}(2p^k\alpha + 1 - \alpha) + \frac{1}{2}\delta\Pi^k(1 + \alpha - 2p^k\alpha) \\ \Leftrightarrow \Pi^k &= \frac{2 - \alpha}{2 - \delta - \delta\alpha + 2p^k\delta\alpha} + \frac{\delta - \delta\alpha + 2p^k\delta\alpha}{2 - \delta - \delta\alpha + 2p^k\delta\alpha}\Pi^{k+1}.\end{aligned}$$

This implies

$$\begin{aligned}\Pi_{\alpha=0}^k &= \frac{2}{2 - \delta} + \frac{\delta}{2 - \delta}\Pi^{k+1} \\ \Pi_{\alpha=1}^k &= \frac{1}{2(1 - \delta + p^k\delta)} + \frac{2p^k\delta}{2(1 - \delta + p^k\delta)}\Pi^{k+1}\end{aligned}$$

¹²As E is indifferent between recommending option 1 and recommending option 2 in case option 2 is his bonus option, his value is as if he always recommended option 1.

where $\Pi_{\alpha=1}^k$ is E's equilibrium value in the supposed equilibrium (where E uses the strategy $\alpha^k = 1$) and $\Pi_{\alpha=0}^k$ is a deviation value that E would obtain if he deviated from the supposed equilibrium strategy by choosing $\alpha^k = 0$ (without changing his strategy for $k' \neq k$). For $\alpha^k = 1$ to be optimal $\Pi_{\alpha=1}^k \geq \Pi_{\alpha=0}^k$ has to hold. However, it is now shown that $\Pi_{\alpha=0}^k > \Pi_{\alpha=1}^k$ for any $\alpha > 0$. This inequality can be written as

$$\begin{aligned} \frac{2}{2-\delta} + \frac{\delta}{2-\delta} \Pi^{k+1} &> \frac{2-\alpha}{2-\delta-\delta\alpha+2p^k\delta\alpha} + \frac{\delta-\delta\alpha+2p^k\delta\alpha}{2-\delta-\delta\alpha+2p^k\delta\alpha} \Pi^{k+1} \\ \Leftrightarrow 4-2\delta-2\delta\alpha+4p^k\alpha\delta+(2-\delta-\delta\alpha+2p^k\alpha\delta)\delta\Pi^{k+1} & \\ &> 4+\alpha\delta-2\delta-2\alpha+(2+\alpha\delta-\delta-2\alpha+4p^k\alpha-2p^k\delta\alpha)\delta\Pi^{k+1} \\ \Leftrightarrow -3\alpha\delta+2\alpha+4\alpha p^k\delta &> (1-\delta)(4p^k\alpha-2\alpha)\delta\Pi^{k+1}. \end{aligned}$$

The latter inequality is true for all $\alpha > 0$ because Π^{k+1} is bounded from above by $1/(1-\delta)$ (which would be E's discounted payoff stream if he always recommended his bonus option and C always continued) and the previous inequality holds with $1/(1-\delta)$ in place of Π^{k+1} :

$$\begin{aligned} -3\alpha\delta+2\alpha+4\alpha p^k\delta &> (1-\delta)(4p^k\alpha-2\alpha)\frac{\delta}{1-\delta} \\ \Leftrightarrow \alpha(2-\delta) &> 0. \end{aligned}$$

This shows that $\alpha^k = 0$ is the only best response to $\beta^k = 1$ and therefore $\Pi_{\alpha=1}^k < \Pi_{\alpha=0}^k$. Consequently, $\beta^k = \alpha^k = 1$ cannot be an equilibrium. \square

Proof of Lemma 8. Proposition 5 implies $V^n = V_O$. Suppose $V^k > V_O$ for some k and let k' be the highest such k . Then $\alpha^{k'}$ must be sufficiently high in order to yield a higher expected payoff than $(1-\delta)V_O$ to C in every period with precision $p^{k'}$. Now consider C's decision problem after a miss in a period with precision $p^{k'}$. As $V^{k'} > V_O$ by the definition of k' , C strictly prefers to continue. Hence, $\beta^{k'} = 1$. However, E's best response to $\beta^{k'} = 1$ is $\alpha^{k'} = 0$, see the proof of Lemma 7. But given that $V^k = V_O$ for all $k > k'$ by the definition of k' and given that $\alpha^{k'} = 0$ clearly $V^{k'} < V_O$ contradicting the definition of k' . Hence, $V^k > V_O$ cannot happen for any k in equilibrium. As C can always guarantee himself a payoff of V_O by ending the game, this concludes the proof. \square

Proof of Proposition 6. As E is mixing in a non-babbling Markov equilibrium when

the bonus option is option 2, his value will equal the value he would get if he always recommended option 1 (keeping C's strategy fixed):

$$\Pi^k = \frac{1}{2} + p^k \delta \Pi^{k+1} + (1 - p^k) \beta^k \delta \Pi^k.$$

Plugging (4.6) in for β^k yields

$$\begin{aligned} \Pi^k &= \frac{1}{2} + p^k \delta \Pi^{k+1} + (1 - p^k) \delta \Pi^{k+1} - \frac{1 - p^k}{2p^k - 1} \\ \Leftrightarrow \Pi^k &= \delta \Pi^{k+1} + \frac{4p^k - 3}{4p^k - 2}. \end{aligned}$$

Recall from the proof of Proposition 5 that $\Pi^n = (4p^n - 3)/[(4p^n - 2)(1 - \delta)]$. Using this as a starting point for backward induction in the previous equation yields (4.8).

Next we will show that Π^k is strictly increasing in k . Let $h(p^k) = (4p^k - 3)/(4p^k - 2)$ and note that $h' > 0$ for $p^k \in (1/2, 1]$. To start, we show by induction that $(1 - \delta)\Pi^k \geq h(p^k)$. This is obviously true for $k = n$. Now suppose $(1 - \delta)\Pi^k \geq h(p^k)$ is true for all $k \geq j + 1$, then $(1 - \delta)\Pi^j = (1 - \delta)\delta \Pi^{j+1} + (1 - \delta)h(p^j) \geq \delta h(p^{j+1}) + (1 - \delta)h(p^j) \geq h(p^j)$ where the first inequality is the induction hypothesis and the second follows from the monotonicity of h . Consequently $(1 - \delta)\Pi^k \geq h(p^k)$ for all $k \in \{1, \dots, n\}$. As $\Pi^{k+1} - \Pi^k = (1 - \delta)\Pi^{k+1} - h(p^k) \geq h(p^{k+1}) - h(p^k) > 0$, it follows that Π^k is strictly increasing in k .

For existence of a non-babbling Markov equilibrium, a $\beta^k \in [0, 1]$ has to exist to make E indifferent between the two recommendations in case option 2 is the bonus option. For $\beta^k = 1$, E strictly prefers to recommend option 2. As the incentives to recommend option 1 are strictly decreasing in β^k , a $\beta^k \in [0, 1]$ will exist if and only if E prefers recommending option 1 (in case option 2 is the bonus option) for $\beta^k = 0$. That is, if

$$\begin{aligned} 1 + (1 - p^k) \delta \Pi^{k+1} &\leq p^k \delta \Pi^{k+1} \\ \Leftrightarrow \Pi^{k+1} &\geq \frac{1}{\delta(2p^k - 1)}. \end{aligned}$$

This condition is most demanding for $k = 1$ because p^k and Π^k are both increasing in k . Hence, a non-babbling Markov equilibrium exists if and only if $\Pi^2 \geq 1/(\delta(2p^1 - 1))$. Plugging in the above derived expression for Π^2 , this is condition (4.7). \square

Proof of Lemma 9. The first claim is proven by a simple strategy copying argument. To show the monotonicity of Π_0^k in k let α_t^k be E's best response strategy to C's m -strategy. More precisely, α_t^k is the probability with which E recommends option 1 when it is not the bonus option (after t misses when the signal precision is p^k). To show that $\Pi_0^{k+1} \geq \Pi_0^k$, we will show that E can achieve a value of Π_0^k at precision $k+1$ (after 0 misses). Note that a signal of precision p^{k+1} is sufficient for a signal of precision p^k . That is, E could inject noise into his signal at precision p^{k+1} in order to end up with a signal of precision p^k . Suppose for all $\tilde{k} \geq k+1$ (and all t) E injects noise into his signal such that the new signal has precision $p^{\tilde{k}-1}$ and then plays the strategy $\hat{\alpha}_t^{\tilde{k}} = \alpha_t^{\tilde{k}-1}$. Equivalently, E can use his improved signal and adjust his behavior to inject some noise in this way.¹³ Clearly, this will yield a value of Π_0^k (at precision $k+1$ after 0 misses). Hence, Π_0^{k+1} has to be at least Π_0^k (and is usually higher as the described strategy is not optimal).

Next we show an intermediate result: $\Pi_t^k \leq \Pi_0^{k+1}$ in every m -equilibrium. To see this, note that E's payoffs are bounded from above by $1/(1-\delta)$, i.e. the value of recommending the bonus option each period and C never stopping the game. Put differently, per period payoffs are below 1. This implies $(1-\delta)\Pi_0^{k+1} \leq 1$. Now suppose, by way of contradiction, $\Pi_t^k > \Pi_0^{k+1}$. Then also $\Pi_{t-1}^k > \Pi_0^{k+1}$ because E can after $t-1$ misses simply recommend his bonus option which would then give him a value of $1 + \delta\Pi_t^k/2 + \delta\Pi_0^{k+1}/2 > 1 + \delta\Pi_0^{k+1} \geq \Pi_0^{k+1}$ where the first inequality uses $\Pi_t^k > \Pi_0^{k+1}$ and the second inequality uses $(1-\delta)\Pi_0^{k+1} \leq 1$. Hence, $\Pi_{t-1}^k > \Pi_0^{k+1}$. Iterating this argument yields $\Pi_0^k > \Pi_0^{k+1}$. However, $\Pi_0^k > \Pi_0^{k+1}$ contradicts the first result of Lemma 9 shown above. Hence, $\Pi_t^k \leq \Pi_0^{k+1}$ holds in every m -equilibrium.

$\Pi_t^k \geq \Pi_{t+1}^k$ is shown using the intermediate result of the previous paragraph. Let E recommend his bonus option after t misses (at precision k). This (possibly non-optimal strategy) yields a value of $1 + \delta\Pi_0^{k+1}/2 + \delta\Pi_{t+1}^k/2 \geq 1 + \delta\Pi_{t+1}^k \geq \Pi_{t+1}^k$ where the first inequality uses $\Pi_{t+1}^k \leq \Pi_0^{k+1}$ (see previous paragraph) and the second inequality uses $\Pi_{t+1}^k \leq 1/(1-\delta)$. As recommending E's bonus option after t misses yields a value of at least Π_{t+1}^k , the result $\Pi_t^k \geq \Pi_{t+1}^k$ follows. \square

Proof of Lemma 10. l^k is the smallest natural number t such that after t consecutive

¹³More precisely, let $\gamma^{\tilde{k}} = (p^{\tilde{k}} - p^{\tilde{k}-1})/(p^{\tilde{k}} - 1/2)$. This is chosen such that drawing from the prior with probability $\gamma^{\tilde{k}}$ and with the counter probability from a signal technology with precision $p^{\tilde{k}}$ yields a signal of precision $p^{\tilde{k}-1}$. If option 1 is the bonus option, let E recommend option 1 with probability $1 - \gamma^{\tilde{k}}/2$ and option 2 with probability $\gamma^{\tilde{k}}/2$. If option 2 is the bonus option, let $\hat{\alpha}_t^{\tilde{k}} = (1 - \gamma^{\tilde{k}}/2)\alpha_t^{\tilde{k}-1} + (\gamma^{\tilde{k}}/2)(1 - \alpha_t^{\tilde{k}-1})$. This yields $\hat{\Pi}_t^{\tilde{k}} = \Pi_t^{\tilde{k}-1}$.

misses, it is optimal for the expert to recommend option 1. We can thus take condition (4.9) and replace $t - 1$ by t . Then, it can be written as $\Pi_{t+1}^k \leq \Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta}$. Hence l^k is the smallest natural number t for which the latter condition holds. More explicitly,

$$\begin{aligned} l^k &= \min\left\{t \in \mathbb{N} \mid \Pi_{t+1}^k \leq \Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta}\right\} \\ &= \max\left\{0, \min\left\{t \in \mathbb{N} \mid \underbrace{\Pi_t^k \leq \Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta}}_{(*)}\right\} - 1\right\}. \end{aligned}$$

Using formula (4.10), we can reformulate the condition (*) via

$$\begin{aligned} (*) &\Leftrightarrow \sum_{j=0}^{m-1-t} (\delta(1-p^k))^j \left(\frac{1}{2} + p^k \delta \Pi_0^{k+1}\right) \leq \Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta} \\ &\Leftrightarrow \sum_{j=0}^{m-1-t} (\delta(1-p^k))^j \leq \underbrace{\frac{\Pi_0^{k+1} - \frac{1}{(2p^k - 1)\delta}}{\frac{1}{2} + p^k \delta \Pi_0^{k+1}}}_{=:P} \\ &\Leftrightarrow \frac{1 - (\delta(1-p^k))^{m-t}}{1 - \delta(1-p^k)} \leq P \\ &\Leftrightarrow \underbrace{1 - (1 - \delta(1-p^k))^{m-t}}_{=:A} \leq \underbrace{(\delta(1-p^k))^{m-t}}_{=:B} \end{aligned}$$

Looking at this last inequality, we see that it is always satisfied if $A \leq 0$ and that it is never satisfied if $P < 0$ (which is equivalent to $\Pi_0^{k+1} < \frac{1}{(2p^k - 1)\delta}$). These cases correspond to $l^k = 0$ and $l^k = m$, respectively. In all the other cases, we can apply the natural logarithm on both sides since they will be positive. We continue:

$$\begin{aligned} &\Rightarrow \ln(A) \leq (m-t) \ln(B) \\ &\Leftrightarrow \frac{\ln(A)}{\ln(B)} \geq m-t \quad (\text{as } \ln(B) < 0) \\ &\Leftrightarrow t \geq m - \frac{\ln(A)}{\ln(B)} \end{aligned}$$

This implies that $l^k = \left\lceil m - 1 - \frac{\ln(A)}{\ln(B)} \right\rceil$ whenever the number inside the ceiling function is larger than -1 and $l^k = 0$ else. \square

Proof of Lemma 11. Define $\tilde{\Pi}_t(\Pi_0)$ (for $t \in \{0, 1, \dots, m\}$) by iterating backwards starting from $\tilde{\Pi}_m = 0$ and using the following formula:

$$\tilde{\Pi}_{t-1}(\Pi_0) = \begin{cases} 1/2 + p^n \delta \Pi_0 + (1 - p^n) \delta \tilde{\Pi}_t(\Pi_0) & \text{if } \Pi_0 - \tilde{\Pi}_t(\Pi_0) \geq \frac{1}{\delta(2p^n - 1)} \\ 1 + \frac{1}{2} \delta \Pi_0 + \frac{1}{2} \delta \tilde{\Pi}_t(\Pi_0) & \text{else.} \end{cases} \quad (17)$$

Note that the case distinction is done such that $\tilde{\Pi}_t$ is continuous ($\tilde{\Pi}_{t-1}(\Pi_0)$ is simply the maximum of the expression in the first and second case). Clearly, the derivative $\tilde{\Pi}'_{m-1}$ exists for almost all values of Π_0 and is in $\{\delta/2, p^n \delta\}$. Hence, $\tilde{\Pi}'_{m-1} < \delta \leq 1$. Iterating backwards, $\tilde{\Pi}_t$ is continuous and its derivative exists for almost all Π_0 . Furthermore, $\tilde{\Pi}'_{t-1}(\Pi_0) \in \{p^n \delta + (1 - p^n) \delta \tilde{\Pi}'_t(\Pi_0), \delta/2 + \delta \tilde{\Pi}'_t(\Pi_0)/2\}$ and therefore – given that $\tilde{\Pi}'_t < \delta$ – we have $\tilde{\Pi}'_{t-1} < \delta$. In particular, $\tilde{\Pi}_0(\Pi_0)$ is continuous and has a derivative (which exists almost everywhere) that is strictly positive and strictly smaller than $\delta \leq 1$. The operator $\tilde{\Pi}_0$ is therefore a contraction and the equation $\tilde{\Pi}_0(\Pi_0) = \Pi_0$ has a unique solution Π_0^* by the contraction mapping theorem.

Next we show that $\Pi_0^* \in (0, 1/(1 - \delta))$. To this purpose it is sufficient to show $\tilde{\Pi}_0(0) > 0$ and $\tilde{\Pi}_0(1/(1 - \delta)) < 1/(1 - \delta)$. Clearly, $\tilde{\Pi}_{t-1}(0) > 0$ and in particular $\tilde{\Pi}_0(0) > 0$ holds. Turning to $\tilde{\Pi}_0(1/(1 - \delta)) < 1/(1 - \delta)$, note that $\tilde{\Pi}_{m-1}(1/(1 - \delta)) < 1/(1 - \delta)$ as both $1/2 + p^n \delta/(1 - \delta) < 1/(1 - \delta)$ and $1 + \delta/(2(1 - \delta)) < 1/(1 - \delta)$ hold. Now proceeding by backward induction $\tilde{\Pi}_{t-1}(1/(1 - \delta)) < 1/(1 - \delta)$ given that $\tilde{\Pi}_t(1/(1 - \delta)) < 1/(1 - \delta)$ as both $1/2 + p^n \delta/(1 - \delta) + (1 - p^n) \delta/(1 - \delta) \leq 1/(1 - \delta)$ and $1 + \delta/(2(1 - \delta)) + \delta/(2(1 - \delta)) \leq 1/(1 - \delta)$ hold. Given that $\Pi_0^* \in (0, 1/(1 - \delta))$, also $\tilde{\Pi}_{t-1}(\Pi_0^*) \in (0, 1/(1 - \delta))$ for all $t \in \{1, \dots, m\}$ by the same steps.

Note that E's value when playing best response against an m -strategy has to satisfy $\Pi_t^n = \tilde{\Pi}_t(\Pi_0^n)$ for all $t \in \{0, \dots, m - 1\}$. As we have just shown, there exists a unique solution to this condition and this solution is feasible, i.e. E's value is in $(0, 1/(1 - \delta))$. E's best response strategy is given by the case distinctions in (17): If $\Pi_0^* - \tilde{\Pi}_t(\Pi_0^*) \geq 1/(\delta(2p^n - 1))$, then E recommends option 1 after $t - 1$ misses. Otherwise, E recommends his bonus option. Finally, note that E's best response is a cutoff strategy as $\tilde{\Pi}_t(\Pi_0^*)$ is decreasing in t . This can be shown as in the proof of Lemma 9. \square

Proof of Proposition 7. By (4.1), an m -equilibrium cannot exist if E always recommends his bonus option. This strategy yields a payoff after 0 misses of

$$\Pi_0 = \sum_{j=0}^{m-1} \left(\frac{\delta}{2}\right)^j (1 + \delta\Pi_0/2) = \frac{(1 + \delta\Pi_0/2)(1 - (\delta/2)^m)}{1 - \delta/2}$$

which can be solved for Π_0 yielding

$$\Pi_0^* = \frac{1 - (\delta/2)^m}{1 - \delta + (\delta/2)^{m+1}}.$$

Always recommending the bonus option is not E's best response if after $m - 1$ misses (4.9) holds with Π_0^* in place of Π_0^{k+1} and zero in place of Π_{t+1}^k , i.e. if

$$\frac{1 - (\delta/2)^m}{1 - \delta + (\delta/2)^{m+1}} \geq \frac{1}{\delta(2p^n - 1)}.$$

If the opposite of this inequality holds, then always recommending the bonus option is E's best response to C's m -strategy (and this best response is unique by Lemma 11) and therefore no m -equilibrium can exist. This gives the condition in (4.13). If (4.13) does not hold, then E's unique best response to C's m -strategy includes recommending option 1 after $m-1$ misses. This implies that C's value when using his m -strategy is strictly above $(1/2)/(1 - \delta)$ (given that V_O satisfies (4.1)) and therefore there exist values of $V_O > (1/2)/(1 - \delta)$ such that C's value is above V_O if C plays an m -strategy and E plays his best response to this strategy. \square

Proof of Proposition 8. We compute E's conditional value depending on the event $w \in \mathcal{W}$ that occurred and then sum over all possible events (making a distinction between histories in which learning level n is reached and those in which the advice relationship was dissolved before). More concretely, we get

$$\Pi_0 = \sum_{\bar{w} \in \mathcal{W}} \mathbb{P}(w = \bar{w}) \mathbb{E}(\Pi_0 | w = \bar{w}).$$

For $\bar{w} \in \mathcal{W}_n$, we get

$$\begin{aligned} \mathbb{P}(w = \bar{w}) &= \prod_{i=1}^{n-1} \mathbb{P}(w_i = \bar{w}_i) = \prod_{i=1}^{n-1} \left(\mathbb{1}_{\{\bar{w}_i \leq l^i\}} \left(\frac{1}{2}\right)^{\bar{w}_i} + \mathbb{1}_{\{\bar{w}_i > l^i\}} \left(\frac{1}{2}\right)^{l^i} (1 - p^i)^{\bar{w}_i - l^i - 1} p^i \right), \\ \mathbb{E}(\Pi_0 | w = \bar{w}) &= \sum_{k=1}^{n-1} H_k \delta^{\sum_{j=1}^{k-1} \bar{w}_j} + \delta^{\sum_{j=1}^{n-1} \bar{w}_j} C_{n-1}. \end{aligned}$$

In the above equation, H_k denotes the expert's expected value in learning level p^k (in which he will spend \bar{w}_k periods). Moreover, C_{n-1} is the expert's continuation value after the $(n-1)$ -th hit at the first period with learning level p^n . Both H_k and C_{n-1} are computed below.

$$H_k = \mathbb{1}_{\{\bar{w}_k \leq l^k\}} \sum_{h=0}^{\bar{w}_k-1} \delta^h + \mathbb{1}_{\{\bar{w}_k > l^k\}} \left(\sum_{h=0}^{l^k-1} \delta^h + \frac{1}{2} \sum_{h=l^k}^{\bar{w}_k-1} \delta^h \right) \quad (18)$$

$$C_{n-1} = \sum_{g=0}^{l^n-1} \left(\frac{\delta}{2}\right)^g + \left(\frac{\delta}{2}\right)^{g+1} C_{n-1} + \sum_{g=l^n}^{m-1} \left(\frac{1}{2}\right)^{l^n} \delta^g (1-p^n)^{g-l^n} p^n C_{n-1} + \delta^g \left(\frac{1}{2}\right)^{g-l^n} \quad (19)$$

$$\Leftrightarrow C_{n-1} = \frac{\sum_{g=0}^{l^n-1} \left(\frac{\delta}{2}\right)^g + \sum_{g=l^n}^{m-1} \delta^g \left(\frac{1}{2}\right)^{g-l^n}}{1 - \sum_{g=0}^{l^n-1} \left(\frac{\delta}{2}\right)^{g+1} - \sum_{g=l^n}^{m-1} \left(\frac{1}{2}\right)^{l^n} \delta^g (1-p^n)^{g-l^n} p^n}$$

In the above computations, (18) follows since the expert will recommend his bonus option l^k times after reaching a new learning level (assuming that all these recommendations produce misses). Only after l^k misses, he will recommend option 1, which yields him $\frac{1}{2}$ per period in expectation, since bonus option and option 1 are drawn independently. Equation (19) reflects the fact that the experts continuation value after $n-1$ hits and after n hits (or more) is the same, since no further learning happens after precision p^n is reached.

For $\bar{w} \in \mathcal{W}_f$, we get

$$\begin{aligned} \mathbb{P}(w = \bar{w}) &= \prod_{i=1}^{len(\bar{w})} \mathbb{P}(w_i = \bar{w}_i) \\ &= \left(\frac{1}{2}\right)^{l^{len(\bar{w})}} (1-p^{len(\bar{w})})^{m-l^{len(\bar{w})}} \prod_{i=1}^{len(\bar{w})-1} \left(\mathbb{1}_{\{\bar{w}_i \leq l^i\}} \left(\frac{1}{2}\right)^{\bar{w}_i} + \mathbb{1}_{\{\bar{w}_i > l^i\}} \left(\frac{1}{2}\right)^{l^i} (1-p^i)^{\bar{w}_i-l^i-1} p^i \right), \end{aligned} \quad (20)$$

$$\mathbb{E}(\Pi_0 | w = \bar{w}) = \sum_{k=1}^{len(\bar{w})-1} H_k \delta^{\sum_{j=1}^{k-1} \bar{w}_j} + \delta^{\sum_{j=1}^{len(\bar{w})-1} \bar{w}_j} \left(\sum_{h=0}^{l^{len(\bar{w})}-1} \delta^h + \frac{1}{2} \sum_{h=l^{len(\bar{w})}}^{m-1} \delta^h \right). \quad (21)$$

Equation (20) follows since m consecutive misses in learning level $len(\bar{w})$ only occur if the bonus option was different from option 1 for $l^{len(\bar{w})}$ periods in a row and the expert failed to generate good advice in the remaining $m - l^{len(\bar{w})}$ periods. Likewise, the term in the brackets in (21) describes the expected payoff of the expert in the

learning level in which m consecutive misses are produced. Putting all the above formulas together yields the desired result. \square

Proof of Proposition 9. The probability of having m consecutive misses conditional on reaching precision level k does of course not only depend on the precision level, but also on the strategy l^k of the expert. More concretely,

$$HR(k) = \left(\frac{1}{2}\right)^{l^k} (1 - p^k)^{m-l^k}.$$

As $1 - p^k < \frac{1}{2}$ for all k by assumption, for (weakly) decreasing values of l (i.e. $l^{k+1} \leq l^k$) the hazard rate $HR(k)$ is (strictly) decreasing in k . When l^k is strictly increasing, then it cannot increase to $l^{k+1} = m$ by Remark 1. Hence, we always have $m - l^{k+1} > 0$ in this case. We now derive the sufficient condition for the hazard rate to be (weakly) decreasing:

$$\begin{aligned} HR(k+1) \leq HR(k) &\Leftrightarrow \left(\frac{1}{2}\right)^{l^{k+1}} (1 - p^{k+1})^{m-l^{k+1}} \leq \left(\frac{1}{2}\right)^{l^k} (1 - p^k)^{m-l^k} \\ &\Leftrightarrow \left(\frac{1}{2(1-p^k)}\right)^{l^{k+1}-l^k} \leq \left(\frac{1-p^k}{1-p^{k+1}}\right)^{m-l^{k+1}} \end{aligned}$$

Since the LHS of the last inequality is increasing in l^{k+1} and decreasing in l^k and the RHS is decreasing in l^{k+1} , it is sufficient to consider $l^{k+1} = m - 1$ and $l^k = 0$ (the extreme cases), which yields

$$\begin{aligned} \left(\frac{1}{2(1-p^k)}\right)^{m-1} \leq \frac{1-p^k}{1-p^{k+1}} &\Leftrightarrow 1-p^{k+1} \leq (1-p^k)^m 2^{m-1} \\ &\Leftrightarrow p^{k+1} \geq 1 - (1-p^k)^m 2^{m-1} \end{aligned}$$

This concludes the proof. \square

Proof of Proposition 10. We will show by induction that $l^k = 1$ implies $l^{k-1} = 1$ in an $m = 1$ equilibrium. However, note first that $l^n = 0$ in an $m = 1$ equilibrium as the consumer would otherwise be better off by ending the advice relationship immediately when reaching precision level p^n .

Now assume that $l^k = 1$ for some $k \in \{2, \dots, n-1\}$. This implies that the expected payoff of the expert when choosing $l^k = 1$, namely $1 + \delta \Pi_0^{k+1}/2$, is greater

or equal than his expected payoff when choosing $l^k = 0$, namely $1/2 + p^k \delta \Pi_0^{k+1}$. Put differently, $1 + \delta \Pi_0^{k+1}/2 \geq 1/2 + p^k \delta \Pi_0^{k+1}$ or equivalently $0 \geq -1/2 + \delta \Pi_0^{k+1}(p^k - 1/2)$. As $0 < \Pi_0^k \leq \Pi_0^{k+1}$ and $1/2 \leq p^{k-1} < p^k$ this inequality implies $0 \geq -1/2 + \delta \Pi_0^k(p^{k-1} - 1/2)$ which is equivalent to saying that the expected payoff of the expert is higher when choosing $l^{k-1} = 1$ than when choosing $l^{k-1} = 0$. Consequently, $l^k = 1$ implies $l^{k-1} = 1$.

Hence, in an $m = 1$ equilibrium $l^k = 1$ for $k \leq \bar{k}$ and $l^k = 0$ for $k > \bar{k}$ for some $\bar{k} \in \{0, \dots, n\}$. The result on V_0^k now readily follows as an increase in k improves the quality of advice in two ways: (i) l^k may decrease and, (ii) p^k increases.

More formally, $V_0^n = p^n(1 + \delta V_0^n) \Leftrightarrow V_0^n = p^n/(1 - p^n \delta)$ and $V_0^k = p^k(1 + \delta V_0^{k+1})$ for $k \in \{\bar{k} + 1, \dots, n - 1\}$. For now let $\bar{k} \leq n - 2$, then $V_0^n > V_0^{n-1}$ holds as $p^n/(1 - \delta p^n) > p^{n-1}(1 + \delta p^n/(1 - \delta p^n)) \Leftrightarrow p^{n-1}/p^n < (1 - \delta p^n)/(1 - \delta p^n) = 1$ which is true by $p^{n-1} < p^n$.¹⁴ Using this as the starting point for backward induction $V_0^k = p^k(1 + \delta V_0^{k+1}) > p^{k-1}(1 + \delta V_0^k) = V_0^{k-1}$ by the induction hypothesis $V_0^{k+1} > V_0^k$ and $p^k > p^{k-1}$ for all $k - 1 > \bar{k}$. The backward induction logic extends to \bar{k} where $V_0^{\bar{k}} = (1 + \delta V_0^{\bar{k}+1})/2 < p^{\bar{k}+1} (1 + \delta V_0^{\bar{k}+2}) = V_0^{\bar{k}+1}$ by $1/2 < p^{\bar{k}+1}$ and $V_0^{\bar{k}+1} < V_0^{\bar{k}+2}$. The backward induction argument continues further for $k < \bar{k}$ as there $V_0^k = (1 + \delta V_0^{k+1})/2 < (1 + \delta V_0^{k+2})/2 = V_0^{k+1}$ where the inequality follows from the induction hypothesis $V_0^{k+1} < V_0^{k+2}$. \square

Proof of Lemma 12. We are using formulas (4.10) and (4.11) to compute the expert's value Π_0^2 for different values of l^2 .

i) $l^2 = 0$

$$\begin{aligned} \Rightarrow \Pi_0^2 &= \frac{1}{2} + p^2 \delta \Pi_0^2 \\ \Leftrightarrow \Pi_0^2 &= \frac{1}{2(1 - p^2 \delta)} \end{aligned}$$

ii) $l^2 = 1$

$$\begin{aligned} \Rightarrow \Pi_0^2 &= 1 + \frac{1}{2} \delta \Pi_0^2 \\ \Leftrightarrow \Pi_0^2 &= \frac{1}{1 - \frac{1}{2} \delta} \end{aligned}$$

¹⁴Clearly, the argument below still holds true for the case $\bar{k} = n - 1$.

This implies that $l^2 = 1$ yields a higher expected value than $l^2 = 0$ for the expert if

$$1 - \frac{1}{2}\delta < 2(1 - p^2\delta)$$

$$\Leftrightarrow \delta < \frac{1}{2p^2 - \frac{1}{2}}$$

holds. The only thing that is left to check now is what the optimal choice for l^1 is in each of the two cases above.

i)

$$l^1 = 0 \Rightarrow \Pi_0^1 = \frac{1}{2} + p^1\delta\Pi_0^2 = \frac{1}{2} + \frac{p^1\delta}{2(1 - p^2\delta)}$$

$$l^1 = 1 \Rightarrow \Pi_0^1 = 1 + \frac{\delta}{2}\Pi_0^2 = 1 + \frac{\delta}{4(1 - p^2\delta)}$$

$$\frac{1}{2} + \frac{p^1\delta}{2(1 - p^2\delta)} > 1 + \frac{\delta}{4(1 - p^2\delta)} \Leftrightarrow \frac{(p^1 - \frac{1}{2})\delta}{2(1 - p^2\delta)} > \frac{1}{2} \Leftrightarrow \delta > \frac{1}{p^1 + p^2 - \frac{1}{2}}$$

Hence, in the case $l^2 = 0$, the expert will choose $l^1 = 0$ if $\delta > \frac{1}{p^1 + p^2 - \frac{1}{2}}$ and he will choose $l^1 = 1$ if $\delta < \frac{1}{p^1 + p^2 - \frac{1}{2}}$ holds.

ii)

$$l^1 = 0 \Rightarrow \Pi_0^1 = \frac{1}{2} + p^1\delta\Pi_0^2 = \frac{1}{2} + \frac{p^1\delta}{1 - \frac{1}{2}\delta}$$

$$l^1 = 1 \Rightarrow \Pi_0^1 = 1 + \frac{\delta}{2}\Pi_0^2 = 1 + \frac{\delta}{2 - \delta}$$

$$\frac{1}{2} + \frac{p^1\delta}{1 - \frac{1}{2}\delta} > 1 + \frac{\delta}{2 - \delta} \Leftrightarrow \delta > \frac{1}{2p^1 - \frac{1}{2}}$$

Since the latter equation is never satisfied for $l^2 = 1$ due to $\delta < \frac{1}{2p^2 - \frac{1}{2}} < \frac{1}{2p^1 - \frac{1}{2}}$, $l^2 = 1$ always implies $l^1 = 1$.

This completes the proof. \square

Proof of Lemma 13. Solving via backward induction, we start determining l^2 by going through three different cases.

1. $l^2 = 0$: Then $\Pi_1^2 = 1/2 + p^2\delta\Pi_0^2$ and $\Pi_0^2 = 1/2 + p^2\delta\Pi_0^2 + (1 - p^2)\delta\Pi_1^2$. Plugging the first expression into the second one and solving for Π_0^2 yields

$$\Pi_0^2 = \frac{(1 + \delta)/2 - \delta p^2/2}{1 - p^2\delta - (1 - p^2)\delta^2 p^2}. \quad (22)$$

2. $l^2 = 1$: Then $\Pi_1^2 = 1/2 + p^2\delta\Pi_0^2$ and $\Pi_0^2 = 1 + \delta\Pi_0^2/2 + \delta\Pi_1^2/2$ which can be solved for

$$\Pi_0^2 = \frac{1 + \delta/4}{1 - \delta/2 - \delta^2 p^2/2}. \quad (23)$$

3. $l^2 = 2$: Then $\Pi_1^2 = 1 + \delta\Pi_0^2/2$ and $\Pi_0^2 = 1 + \delta\Pi_0^2/2 + \delta\Pi_1^2/2$ which can be solved for

$$\Pi_0^2 = \frac{1 + \delta/2}{1 - \delta/2 - \delta^2/4}. \quad (24)$$

Therefore, $l^2 = 1$ is the expert's best response if and only if

$$\frac{1 + \delta/4}{1 - \delta/2 - \delta^2 p^2/2} > \max \left\{ \frac{1 + \delta/2}{1 - \delta/2 - \delta^2/4}, \frac{(1 + \delta)/2 - \delta^2 p^2/2}{1 - p^2\delta - (1 - p^2)\delta^2 p^2} \right\}.$$

Conditional on $l^2 = 1$ being the expert's best response in learning level 2, we will now check under which conditions $l^1 = 0$ is the expert's best response in learning level 1. Again, we have to go through three cases.

1. $l^1 = 0$: Then $\Pi_1^1 = 1/2 + p^1\delta\Pi_0^2$ and $\Pi_0^1 = 1/2 + p^1\delta\Pi_0^2 + (1 - p^1)\delta\Pi_1^1$. Plugging in yields

$$\Pi_0^1 = \frac{1}{2} + \frac{(1 - p^1)\delta}{2} + \frac{p^1\delta(1 + (1 - p^1)\delta)(1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2}.$$

2. $l^1 = 1$: Then $\Pi_1^1 = 1/2 + p^1\delta\Pi_0^2$ and $\Pi_0^1 = 1 + \delta\Pi_0^2/2 + \delta\Pi_1^1/2$. Plugging in yields

$$\Pi_0^1 = 1 + \delta/4 + \frac{(\delta/2 + \delta^2 p^1/2)(1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2}.$$

3. $l^1 = 2$: Then $\Pi_1^1 = 1 + \delta\Pi_0^2/2$ and $\Pi_0^1 = 1 + \delta\Pi_0^2/2 + \delta\Pi_1^1/2$. Plugging in yields

$$\Pi_0^1 = 1 + \delta/2 + \frac{(\delta/2 + \delta^2/4)(1 + \delta/4)}{1 - \delta/2 - \delta^2 p^2/2}.$$

Therefore, $l^1 = 0$ will be the expert's best response if and only if

$$\frac{1}{2} + \frac{(1-p^1)\delta}{2} + \frac{p^1\delta(1+(1-p^1)\delta)(1+\delta/4)}{1-\delta/2-\delta^2p^2/2}$$

$$\geq \max \left\{ 1 + \delta/4 + \frac{(\delta/2 + \delta^2p^1/2)(1 + \delta/4)}{1 - \delta/2 - \delta^2p^2/2}, 1 + \delta/2 + \frac{(\delta/2 + \delta^2/4)(1 + \delta/4)}{1 - \delta/2 - \delta^2p^2/2} \right\}.$$

□

B Experiment Instructions

B.1 Control Group

Freiwilligkeit des Experimentes

Die Teilnahme an diesem Experiment ist freiwillig. Sie können die Teilnahme jederzeit ohne Angabe von Gründen abbrechen.

Instruktionen

Bitte lesen Sie die folgenden Instruktionen sorgfältig. Vor dem Experiment bekommen Sie einige Kontrollfragen gestellt und Sie können bei korrekter Beantwortung Geld gewinnen. Konkret werden Ihnen **vier Kontrollfragen** gestellt. Hiervon wird nach Ihren Antworten eine zufällig ausgewählt und wenn Ihre Antwort auf diese Frage beim ersten Versuch richtig war, bekommen Sie eine zusätzliche Auszahlung von **1,00€**.

Im Folgenden werden Sie zufällig in Zweiergruppen eingeteilt und werden mit Ihrem zugeteilten Spielpartner ein Spiel spielen. In diesem Spiel können Sie Spielpunkte erspielen. Auf Basis dieser Spielpunkte wird am Ende Ihre Auszahlung ermittelt, was weiter unten erläutert wird. Zusätzlich erhalten Sie eine hiervon unabhängige Auszahlung von **4,00€** für das Erscheinen und Ihre Teilnahme am Experiment. In dem Spiel werden Sie zufällig entweder die Rolle von Spieler A oder von Spieler B übernehmen. Das Spiel wird nun beschrieben und danach anhand eines Beispiels für zwei Spielrunden veranschaulicht. Dort sehen Sie auch beispielhaft die Bildschirmanzeigen, die beiden Spielern jeweils angezeigt werden.

Entscheidungen der Spieler

Das Spiel wird über mehrere Runden gespielt und in jeder Runde hat Spieler A die Wahl zwischen Option 1 und Option 2 und Spieler B entscheidet in der Folge, ob eine weitere Runde des Spiels gespielt wird. Es gibt in jeder Runde vier mögliche Fälle, die alle gleich wahrscheinlich sind und vor jeder neuen Runde zufällig bestimmt werden.

Fälle	Auszahlungen der Spieler bei Wahl von	
	Option 1	Option 2
1. Fall	A: 1 Punkt, B: 1 Punkt	A: 0 Punkte, B: 0 Punkte
2. Fall	A: 0 Punkte, B: 0 Punkte	A: 1 Punkt, B: 1 Punkt
3. Fall	A: 1 Punkt, B: 0 Punkte	A: 0 Punkte, B: 1 Punkt
4. Fall	A: 0 Punkte, B: 1 Punkt	A: 1 Punkt, B: 0 Punkte

Abbildung 1: Übersicht über die möglichen Auszahlungen für Spieler A und B

In jedem möglichen Fall erhält also jeder Spieler eine Auszahlung von 1 von **genau einer** der beiden Optionen, während die andere Option ihm eine Auszahlung von 0 gibt. Die Option mit der höheren Auszahlung kann entweder für beide Spieler die gleiche oder aber eine unterschiedliche sein.

In jeder Spielrunde tritt genau einer der obigen Fälle ein, aber keiner der Spieler weiß mit Sicherheit, welcher das ist. Spieler A bekommt jedoch immer angezeigt für welche der Optionen er einen Punkt erhält. Darüber hinaus erhält er einen automatisch erzeugten *Hinweis* darüber, welche Option Spieler B einen Punkt einbringen *könnte*. Dieser Hinweis ist immer mit einer **Wahrscheinlichkeit von 82% korrekt** und mit einer **Wahrscheinlichkeit von 18% inkorrekt**.

Nach der Entscheidung von Spieler A werden beide Spieler über ihre daraus resultierenden Auszahlungen informiert. Spieler B erfährt hierbei nur, ob er eine Auszahlung von 1 oder 0 (Spielpunkten) erhält und nicht, was der Hinweis von Spieler A war oder welche Auszahlung Spieler A erhalten hat. Spieler A wird hingegen auch über die Auszahlung von Spieler B informiert. Spieler B kann also keine der Optionen selbst wählen, sondern erhält seine Auszahlung abhängig von der Wahl von Spieler A. Im Anschluss daran kann Spieler B entscheiden, ob er das Spiel *beenden* oder für eine weitere Runde *fortführen* möchte.

- Spieler B wählt *fortführen*:


In diesem Fall wird mit einer Wahrscheinlichkeit von 90% eine weitere Runde des Spiels gespielt.

Mit einer Wahrscheinlichkeit von 10% **endet das Spiel** trotz der Entscheidung von Spieler B das Spiel fortzuführen (sonst könnte das Spiel theoretisch unendlich lange dauern). Beide Spieler erhalten ihre bis dahin erspielten Spielpunkte. Beide Spieler erhalten die Nachricht, dass das Spiel *exogen* beendet wurde.

- Spieler B wählt *beenden*:

In diesem Fall bekommt Spieler B zusätzlich **5 Spielpunkte** gutgeschrieben, Spieler A erhält keine weiteren Punkte. Das Spiel ist zu Ende und beide Spieler werden darüber informiert, dass Spieler B das Spiel beendet hat.

Jede Runde des Spiels kann wie folgt in einem Schaubild veranschaulicht werden:



Neues Spiel mit neuem Spielpartner

Sobald ein Spiel für alle Spieler beendet ist (entweder exogen oder weil alle Spieler B ihr Spiel beendet haben), werden die Spielpartner neu zugelost. Jeder behält hierbei jedoch seine Rolle als Spieler A oder Spieler B und bekommt zufällig einen Spieler des anderen Typs zugelost. Das Spiel wird erneut gestartet. Insgesamt werden **10 Spiele** mit wechselnden Spielpartnern durchgeführt. Am Ende wird zufällig **eines** der 10 Spiele ausgewählt und die dort erspielte Punktzahl wird nach Beendigung des Experimentes (zusammen mit der festen Auszahlung) ausgezahlt. Ein Spielpunkt entspricht hierbei **1,00€**.

Beispiel

In dem folgenden Beispiel (siehe Abbildung 1) wählt Spieler A in der ersten Runde Option 1 (oben links im Bild). Im Folgenden werden beide Spieler darüber informiert, dass diese Wahl Spieler B eine Auszahlung von 0 einbringt (zu sehen ist nur der Bildschirm von Spieler B, oben rechts). Spieler A kann so feststellen, dass sein Hinweis über Spieler B in Runde 1 korrekt war, da der Hinweis Option 2 lautete und Option 1 Spieler B eine Auszahlung von 0 einbrachte. Somit hätte

Option 2 tatsächlich in einer Auszahlung von 1 für Spieler B resultiert. Spieler B weiß allerdings weder welchen Hinweis Spieler A erhalten hat, noch ob Spieler A diesem Hinweis gefolgt ist.

<p>Erzielte Auszahlungen für Spieler B</p> <p>Sie erhalten eine Auszahlung von einem Punkt für Option 1.</p> <p>Der Hinweis für Spieler B lautet: Option 2. <small>Erläuterung: Wenn Sie diese Option wählen, bekommt Spieler A also mit einer Wahrscheinlichkeit von 0,82 eine Auszahlung von 1 und sonst eine Auszahlung von 0. Wenn Sie die andere Option wählen, bekommt Spieler B mit einer Wahrscheinlichkeit von 0,82 eine Auszahlung von 1 und sonst eine Auszahlung von 0.</small></p> <p>Welche Option möchten Sie wählen?</p> <input type="text" value="1"/> <p>Weiter</p>	<p>Spieler A hat die Option ausgewählt, die Ihnen eine Auszahlung von 0 einbringt.</p> <p>Erzielte Auszahlungen 0.</p> <p>Möchten Sie das Spiel beenden oder fortführen?</p> <p><input type="radio"/> Spiel beenden <input type="radio"/> Spiel fortführen</p> <p>Weiter</p>
<p>Erzielte Auszahlungen für Spieler B 0,</p> <p>Sie erhalten eine Auszahlung von einem Punkt für Option 2.</p> <p>Der Hinweis für Spieler B lautet: Option 2. <small>Erläuterung: Wenn Sie diese Option wählen, bekommt Spieler A also mit einer Wahrscheinlichkeit von 0,82 eine Auszahlung von 1 und sonst eine Auszahlung von 0. Wenn Sie die andere Option wählen, bekommt Spieler B mit einer Wahrscheinlichkeit von 0,82 eine Auszahlung von 1 und sonst eine Auszahlung von 0.</small></p> <p>Welche Option möchten Sie wählen?</p> <input type="text" value="2"/> <p>Weiter</p>	<p>Spieler A hat die Option ausgewählt, die Ihnen eine Auszahlung von 1 einbringt.</p> <p>Erzielte Auszahlungen 0,1.</p> <p>Möchten Sie das Spiel beenden oder fortführen?</p> <p><input type="radio"/> Spiel beenden <input type="radio"/> Spiel fortführen</p> <p>Weiter</p>

Abbildung 2: Ein Beispiel für die ersten zwei Spielrunden

Im Beispiel entscheidet sich Spieler B für "Spiel fortführen" und es wird eine zweite Runde gespielt. Nun entscheidet sich Spieler A für Option 2 (Bild unten links). Diese Wahl führt zu einer Auszahlung von 1 für Spieler B (siehe Bild unten rechts). Spieler B kann nun wieder entscheiden, ob er das Spiel fortführen oder beenden möchte.

Ende des Experimentes

Zum Ende des Experimentes bekommen Sie noch ein paar Fragen gestellt, bei denen Sie teilweise Geld gewinnen können (dies ist dann jeweils vor Beantwortung der Fragen erklärt). Zuletzt geben Sie über ein Formular Ihre Auszahlungsdaten ein, die von der Universität zur Tätigung der Zahlung benötigt werden.

B.2 Treatment Group

Freiwilligkeit des Experimentes

Die Teilnahme an diesem Experiment ist freiwillig. Sie können die Teilnahme jederzeit ohne Angabe von Gründen abbrechen.

Instruktionen

Bitte lesen Sie die folgenden Instruktionen sorgfältig. Vor dem Experiment bekommen Sie einige Kontrollfragen gestellt und Sie können bei korrekter Beantwortung Geld gewinnen. Konkret werden Ihnen **fünf Kontrollfragen** gestellt. Hiervon wird nach Ihren Antworten eine zufällig ausgewählt und wenn Ihre Antwort auf diese Frage beim ersten Versuch richtig war, bekommen Sie eine zusätzliche Auszahlung von **1,00€**.

Im Folgenden werden Sie zufällig in Zweiergruppen eingeteilt und werden mit Ihrem zugeteilten Spielpartner ein Spiel spielen. In diesem Spiel können Sie Spielpunkte erspielen. Auf Basis dieser Spielpunkte wird am Ende Ihre Auszahlung ermittelt, was weiter unten erläutert wird. Zusätzlich erhalten Sie eine hiervon unabhängige Auszahlung von **4,00€** für das Erscheinen und Ihre Teilnahme am Experiment. In dem Spiel werden Sie zufällig entweder die Rolle von Spieler A oder von Spieler B übernehmen. Das Spiel wird nun beschrieben und danach anhand eines Beispiels für zwei Spielrunden veranschaulicht. Dort sehen Sie auch beispielhaft die Bildschirmanzeigen, die beiden Spielern jeweils angezeigt werden.

Entscheidungen der Spieler

Das Spiel wird über mehrere Runden gespielt und in jeder Runde hat Spieler A die Wahl zwischen Option 1 und Option 2 und Spieler B entscheidet in der Folge, ob eine weitere Runde des Spiels gespielt wird. Es gibt in jeder Runde vier mögliche Fälle, die alle gleich wahrscheinlich sind und vor jeder neuen Runde zufällig bestimmt werden.

Fälle	Auszahlungen der Spieler bei Wahl von	
	Option 1	Option 2
1. Fall	A: 1 Punkt, B: 1 Punkt	A: 0 Punkte, B: 0 Punkte
2. Fall	A: 0 Punkte, B: 0 Punkte	A: 1 Punkt, B: 1 Punkt
3. Fall	A: 1 Punkt, B: 0 Punkte	A: 0 Punkte, B: 1 Punkt
4. Fall	A: 0 Punkte, B: 1 Punkt	A: 1 Punkt, B: 0 Punkte

Abbildung 1: Übersicht über die möglichen Auszahlungen für Spieler A und B

In jedem möglichen Fall erhält also jeder Spieler eine Auszahlung von 1 von **genau einer** der beiden Optionen, während die andere Option ihm eine Auszahlung

von 0 gibt. Die Option mit der höheren Auszahlung kann entweder für beide Spieler die gleiche oder aber eine unterschiedliche sein.

In jeder Spielrunde tritt genau einer der obigen Fälle ein, aber keiner der Spieler weiß mit Sicherheit, welcher das ist. Spieler A bekommt jedoch immer angezeigt für welche der Optionen er einen Punkt erhält. Darüber hinaus erhält er einen automatisch erzeugten *Hinweis* darüber, welche Option Spieler B einen Punkt einbringen *könnte*. Dieser Hinweis ist in der ersten Runde mit einer **Wahrscheinlichkeit von 82% korrekt** und mit einer **Wahrscheinlichkeit von 18% inkorrekt**. Die Wahrscheinlichkeit, mit der der Hinweis korrekt ist, nennen wir in dem Experiment die **Hinweisstärke**. Sie wird immer als Dezimalzahl angegeben. Eine Hinweisstärke von 0,82 entspricht zum Beispiel einer Wahrscheinlichkeit von 82%, eine Hinweisstärke von 0,84 entspricht 84%, usw.

Nach der Entscheidung von Spieler A werden beide Spieler über ihre daraus resultierenden Auszahlungen informiert. Spieler B erfährt hierbei nur, ob er eine Auszahlung von 1 oder 0 (Spielpunkten) erhält und nicht, was der Hinweis von Spieler A war oder welche Auszahlung Spieler A erhalten hat. Spieler A wird hingegen auch über die Auszahlung von Spieler B informiert. Spieler B kann also keine der Optionen selbst wählen, sondern erhält seine Auszahlung abhängig von der Wahl von Spieler A. Im Anschluss daran kann Spieler B entscheiden, ob er das Spiel *beenden* oder für eine weitere Runde *fortführen* möchte.

- Spieler B wählt *fortführen*:

In diesem Fall wird mit einer Wahrscheinlichkeit von 90% eine weitere Runde des Spiels gespielt. Falls Spieler B in der aktuellen Runde eine Auszahlung von einem Spielpunkt erhalten hat, wird in den folgenden Runden der Hinweis, den Spieler A erhält, *verbessert*: Die Wahrscheinlichkeit, mit der der Hinweis korrekt ist, **erhöht sich um 2%** (die Hinweisstärke erhöht sich also um 0,02). Falls Spieler B in der aktuellen Runde eine Auszahlung von null Spielpunkten erhalten hat, bleibt die Hinweisstärke genau wie in der vorherigen Runde.


Mit einer Wahrscheinlichkeit von 10% **endet das Spiel** trotz der Entscheidung von Spieler B das Spiel fortzuführen (sonst könnte das Spiel theoretisch unendlich lange dauern). Beide Spieler erhalten ihre bis dahin erspielten Spielpunkte. Beide Spieler erhalten die Nachricht, dass das Spiel *exogen* beendet wurde.

- Spieler B wählt *beenden*:

In diesem Fall bekommt Spieler B zusätzlich **5 Spielpunkte** gutgeschrieben, Spieler A erhält keine weiteren Punkte. Das Spiel ist zu Ende und beide Spieler werden darüber informiert, dass Spieler B das Spiel beendet hat.

Sofern das Spiel über mehrere Runden fortgeführt wird, verbessert sich der Hinweis auch in folgenden Runden (sofern Spieler B eine Auszahlung von einem Punkt erhält). Hierbei erhöht sich die Wahrscheinlichkeit, dass der Hinweis definitiv korrekt ist jeweils um **2%**. Die maximale Wahrscheinlichkeit ist jedoch **90%**. Sollte in einer Runde also diese Wahrscheinlichkeit erreicht sein und Spieler B erhält in dieser Runde nochmals eine Auszahlung von 1, so bleibt die Wahrscheinlichkeit auch in allen folgenden Runden bei 90%.

Jede Runde des Spiels kann wie folgt in einem Schaubild veranschaulicht werden:



Neues Spiel mit neuem Spielpartner

Sobald ein Spiel für alle Spieler beendet ist (entweder exogen oder weil alle Spieler B ihr Spiel beendet haben), werden die Spielpartner neu zugelost. Jeder behält hierbei jedoch seine Rolle als Spieler A oder Spieler B und bekommt zufällig einen Spieler des anderen Typs zugelost. Das Spiel wird erneut gestartet. Insgesamt werden **10 Spiele** mit wechselnden Spielpartnern durchgeführt. Am Ende wird zufällig **eines** der 10 Spiele ausgewählt und die dort erspielte Punktzahl wird nach Beendigung des Experimentes (zusammen mit der festen Auszahlung) ausgezahlt. Ein Spielpunkt entspricht hierbei **1,00€**.

Beispiel

In dem folgenden Beispiel (siehe Abbildung 1) wählt Spieler A in der ersten Runde Option 1 (oben links im Bild). Im Folgenden werden beide Spieler darüber informiert, dass diese Wahl Spieler B eine Auszahlung von 0 einbringt (zu sehen ist nur der Bildschirm von Spieler B, oben rechts). Spieler A kann so feststellen, dass sein Hinweis über Spieler B in Runde 1 korrekt war, da der Hinweis Option

2 lautete und Option 1 Spieler B eine Auszahlung von 0 einbrachte. Somit hätte Option 2 tatsächlich in einer Auszahlung von 1 für Spieler B resultiert. Spieler B weiß allerdings weder welchen Hinweis Spieler A erhalten hat, noch ob Spieler A diesem Hinweis gefolgt ist.

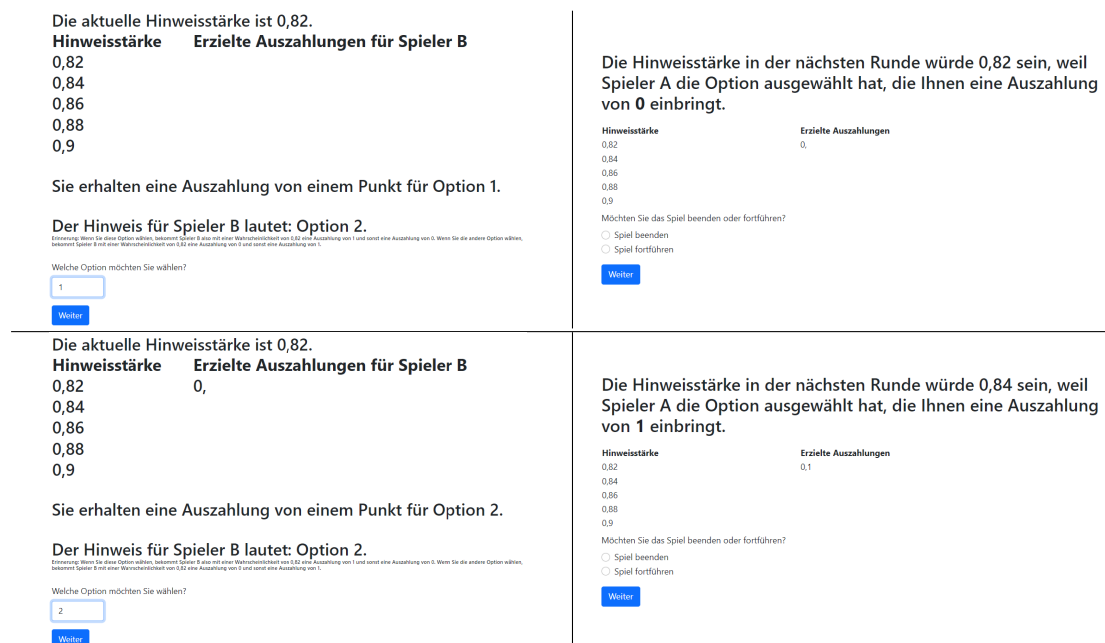


Abbildung 2: Ein Beispiel für die ersten zwei Spielrunden

Im Beispiel entscheidet sich Spieler B für "Spiel fortführen" und es wird eine zweite Runde gespielt. In der zweiten Runde ist die Hinweisstärke dann wiederum 0,82, da Spieler B in der ersten Runde eine Auszahlung von 0 erreicht hat. Nun entscheidet sich Spieler A für Option 2 (Bild unten links). Diese Wahl führt zu einer Auszahlung von 1 für Spieler B (siehe Bild unten rechts). Spieler B kann nun wieder entscheiden, ob er das Spiel fortführen oder beenden möchte und wird darüber informiert, dass die Hinweisstärke in der nächsten Runde 0,84 wäre. Hätte Spieler B in der zweiten Runde eine Auszahlung von 0 erhalten, so wäre die Hinweisstärke in der nächsten Runde weiterhin bei 0,82 geblieben. Die Hinweisstärke erhöht sich immer nur dann, wenn Spieler B in einer Runde eine Auszahlung von 1 erhält.

Ende des Experimentes

Zum Ende des Experimentes bekommen Sie noch ein paar Fragen gestellt, bei denen Sie teilweise Geld gewinnen können (dies ist dann jeweils vor Beantwortung

der Fragen erklärt). Zuletzt geben Sie über ein Formular Ihre Auszahlungsdaten ein, die von der Universität zur Tatigung der Zahlung benotigt werden.

C Additional Results and Robustness Checks

In this section, we will provide additional results and robustness checks related to the experimental results given in Section 4.6. We will have a further look at advice quality, welfare distribution and hazard rates in turn.

C.1 Advice Quality

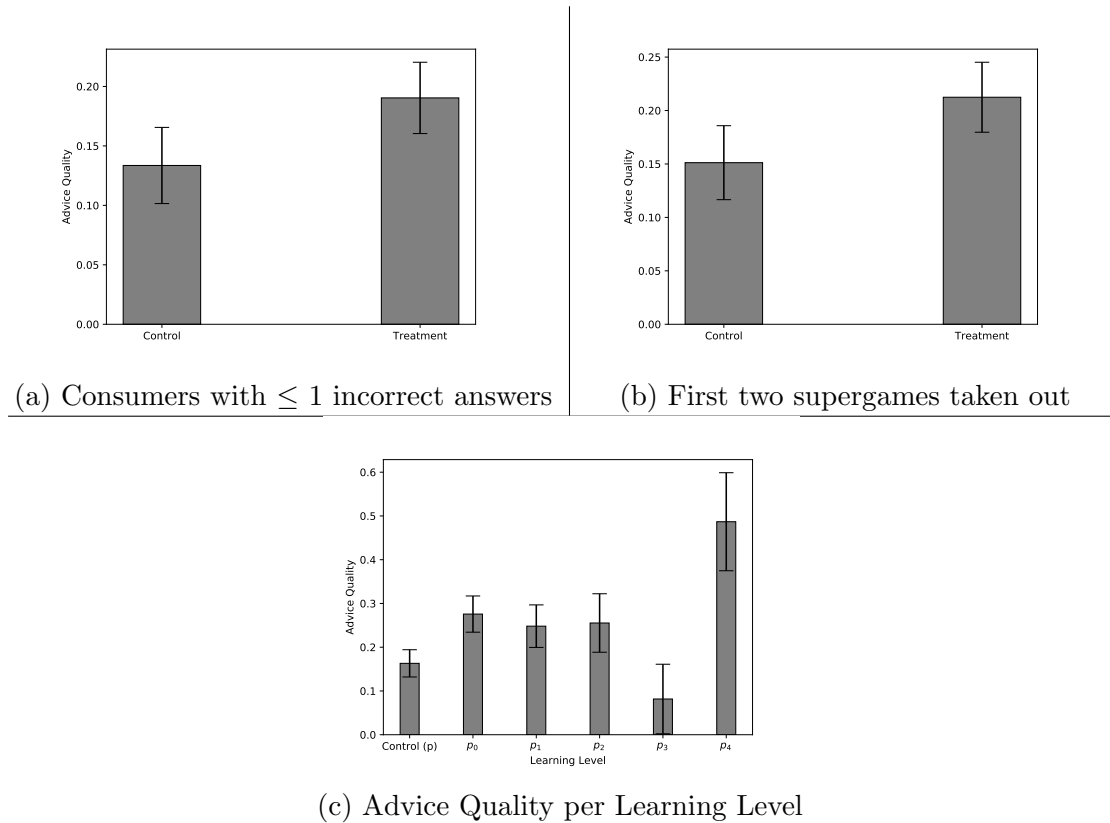


Figure 7: Robustness Checks for Advice Quality

Figure 7 shows the robustness checks for advice quality. In Figure 7a, we only considered those experts who gave at most one incorrect answer to the check questions. We can see that the advice quality in the treatment group is still higher, but the difference becomes a bit less significant. The same happens when we take out the first two supergames in each session, where players might still have been learning the game. This is shown in Figure 7b. Lastly, we looked at advice quality

per learning level in the treatment group. The results can be seen in Figure 7c. It turns out that advice in the learning levels p_0, p_1, p_2 and p_4 is significantly higher than that given in the control group. However, the average advice quality in learning level p_3 is lower than in the control group. A potential explanation for the low advice quality in this level is the gambling effect: Experts feel that their signal strength is sufficiently high to generate fitting advice on the spot such that they will take their bonus and hope to appease the consumer in the next period. It is also noteworthy that the advice quality in learning level p_4 is significantly higher than in all other learning levels as well as in the control group. This effect cannot be explained by a better signal quality, since advice quality is measured by the share of tradeoff-situations (bonus option = option 2), in which the adviser decides to give useful advice instead of receiving his bonus. The signal quality only affects how often this decision will actually translate to the intended payoff of one to the consumer. A reason for the high advice quality in learning level p_4 could be a selection effect: The majority of advisers who reached learning level p_4 in their advice relationship probably did so because they gave good advice in the past and they might have an intrinsic motivation to give good advice and/or value long-lasting relationships a lot. Another explanation could be reciprocity: Advisers reward consumers for their loyalty over the last rounds by giving better advice.

Overall, we conclude that the difference in advice quality between control group and treatment group seems to be quite robust.

C.2 Welfare Analysis

The results of the robustness checks for consumer welfare can be seen in Figure 8. Overall, the observation that consumer welfare does not significantly differ between control and treatment group is very robust. When we take out the supergames with less than three rounds (Figure 8a) or the first two supergames of each session (Figure 8c) or those consumers with two or more incorrect answers to check questions (Figure 8b), there is no significant difference in consumer welfare between control and treatment group. We also had a look at total welfare, the sum of consumer and expert payoffs. As can be seen in Figure 8d, there is no significant difference between control and treatment group, either. This also implies that expert payoffs in control and treatment group are not significantly different from each other.

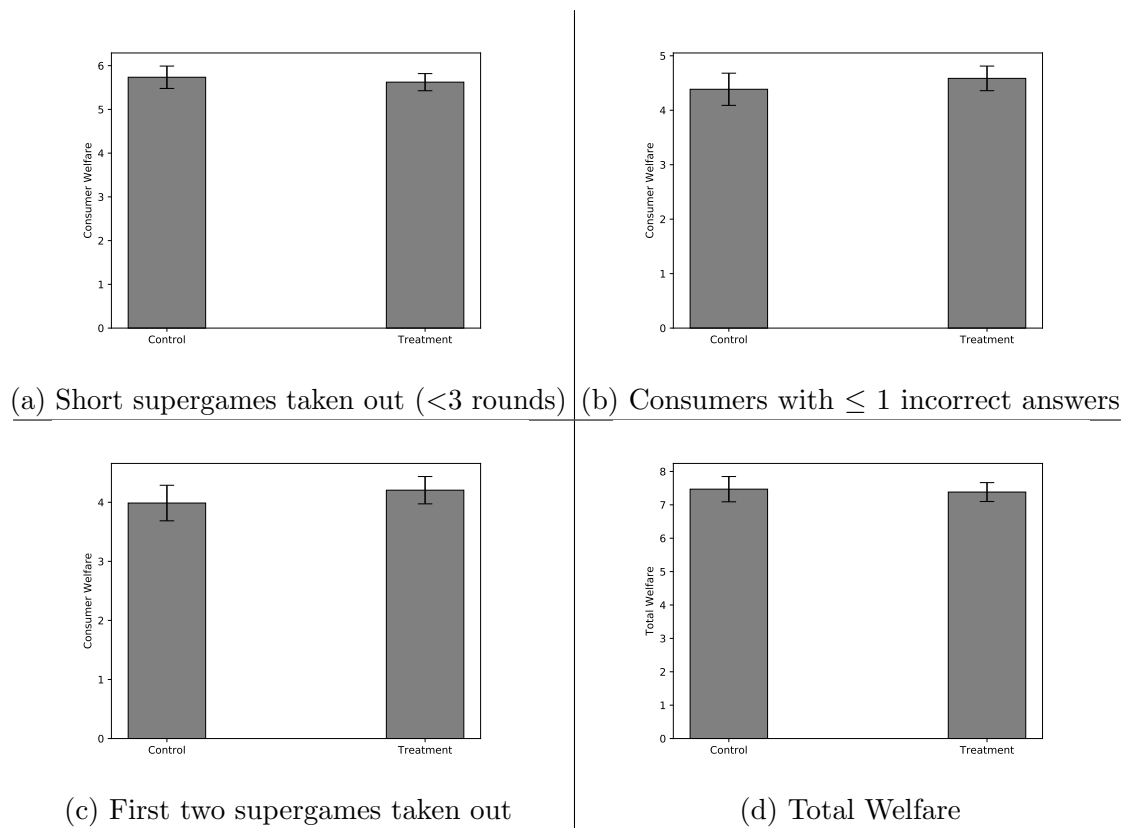




Figure 8: Robustness Checks for Consumer Welfare

C.3 Hazard Rates

Figure 9 shows the robustness checks we performed for the hazard rates. Our finding that hazard rates are significantly lower in learning levels p_0 and p_1 proves to be robust. Both excluding consumers with two or more incorrect answers to check questions (Figure 9a) and taking out the first two supergames of each session (Figure 9b) leads to a shape very similar to the one in Figure 4.4.



(a) Consumers with ≤ 1 incorrect answers



(b) First two supergames taken out

Figure 9: Robustness Checks for Hazard Rates

References

- Acquisti, A., L. Brandimarte, and G. Loewenstein (2015). Privacy and human behavior in the age of information. *Science* 347(6221), 509–514.
- Akbarpour, M., S. Li, and S. O. Gharan (2020). Thickness and Information in Dynamic Matching Markets. *Journal of Political Economy* 128(3), 783–815.
- Anderson, A., J. Rosen, J. Rust, and K.-P. Wong (2021). Disequilibrium Play in Tennis. Working paper, Georgetown University, Department of Economics.
- Arnosti, N., R. Johari, and Y. Kanoria (2014). Managing congestion in decentralized matching markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 451–451.
- Baccara, M., S. Lee, and L. Yariv (2020). Optimal dynamic matching. *Theoretical Economics* 15(3), 1221–1278.
- Becker, R. A., S. K. Chakrabarti, W. Geller, B. Kitchens, and M. Misiurewicz (2007). Dynamics of the Nash map in the game of Matching Pennies. *Journal of Difference Equations and Applications* 13(2-3), 223–235.
- Benabou, R. and G. Laroque (1992). Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. *Quarterly Journal of Economics* 107(3), 921–958.
- Blume, A., E. K. Lai, and W. Lim (2020). Strategic information transmission: A survey of experiments and theoretical foundations. In *Handbook of experimental game theory*, pp. 311–347. Edward Elgar Publishing.
- Borel, E. (1921). La théorie du jeu et les équations intégrales à noyau symétrique gauche. *Comptes Rendus de l'Académie des Sciences* 173, 1304–08. English translation by Savage, L.: The Theory of Play and Integral Equations with Skew Symmetric Kernels. *Econometrica* 21, 97–100 (1953).

- Byers, S., L. F. Cranor, D. Kormann, and P. McDaniel (2004). Searching for Privacy: Design and Implementation of a P3P-Enabled Search Engine. In *International Workshop on Privacy Enhancing Technologies*, pp. 314–328. Springer.
- Cardoso, J. and P. Diniz (2009). Defending Against Terrorism, Natural Disaster, and All Hazards. In *Game Theoretic Risk Analysis of Security Threats*, pp. 65–97. Springer.
- Çetin, G. S., W. Dai, Y. Doröz, W. J. Martin, and B. Sunar (2016). Blind Web Search: How far are we from a privacy preserving search engine? *IACR Cryptology ePrint Archive 2016*, 801.
- Chade, H., G. Lewis, and L. Smith (2014). Student Portfolios and the College Admissions Problem. *Review of Economic Studies* 81(3), 971–1002.
- Che, Y.-K. and Y. Koh (2016). Decentralized College Admissions. *Journal of Political Economy* 124(5), 1295–1338.
- Chellappa, R. K. and R. G. Sin (2005). Personalization versus Privacy: An Empirical Examination of the Online Consumer’s Dilemma. *Information technology and management* 6(2), 181–202.
- Chen, D. L., M. Schonger, and C. Wickens (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88–97.
- Cheng, L. and T. Yu (2018). Nash Equilibrium-Based Asymptotic Stability Analysis of Multi-Group Asymmetric Evolutionary Games in Typical Scenario of Electricity Market. *IEEE Access* 6, 32064–32086.
- Coles, P., J. Cawley, P. B. Levine, M. Niederle, A. E. Roth, and J. J. Siegfried (2010). The Job Market for New Economists: A Market Design Perspective. *Journal of Economic Perspectives* 24(4), 187–206.
- Crawford, V. P. and J. Sobel (1982). Strategic Information Transmission. *Econometrica* 50(6), 1431–1451.
- Edelman, B., M. Ostrovsky, and M. Schwarz (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97(1), 242–259.

- Edelman, B. and M. Schwarz (2010). Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions. *American Economic Review* 100(2), 597–602.
- Eliasz, K. and R. Spiegler (2011). A Simple Model of Search Engine Pricing. *Economic Journal* 121(556), 329–339.
- Friedman, D. and B. Sinervo (2016). *Evolutionary Games in Natural, Social, and Virtual Worlds*. Oxford University Press.
- Friedman, M. and R. Friedman (1980). *Free to Choose: A Personal Statement*. Harcourt Brace Jovanovich.
- Fudenberg, D. and D. K. Levine (1998). *The Theory of Learning in Games*, Volume 2. MIT Press.
- Gentle, J. E. (2009). *Computational Statistics*. Springer.
- Golman, R. and S. E. Page (2009). General Blotto: Games of Allocative Strategic Mismatch. *Public Choice* 138(3-4), 279–299.
- Gong, L., J. Gao, and M. Cao (2018). Evolutionary Game Dynamics for Two Interacting Populations in a Co-evolving Environment. In *2018 IEEE Conference on Decision and Control (CDC)*, pp. 3535–3540. IEEE.
- Gramb, M. (2022). Game Preparation and Experience. *Available at SSRN: <http://dx.doi.org/10.2139/ssrn.4150612>*.
- Gramb, M. and C. Schottmüller (2022a). Anonymous or personal? A simple model of repeated personalized advice. *University of Cologne, Working Paper Series in Economics No. 105*.
- Gramb, M. and C. Schottmüller (2022b). Anonymous or personal? Repeated personalized advice in the lab. *AEA RCT Registry*. January 13. doi:10.1257/rct.8682.
- Gramb, M. and J. Teichgräber (2023). Congestion and Market Thickness in Decentralized Matching Markets. *Available at SSRN: <http://dx.doi.org/10.2139/ssrn.4333144>*.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association* 1(1), 114–125.

- He, Y. and T. Magnac (2022). Application Costs and Congestion in Matching Markets. *The Economic Journal* 132(648), 2918–2950.
- Hernández, D. G. and D. H. Zanette (2013). Evolutionary Dynamics of Resource Allocation in the Colonel Blotto Game. *Journal of Statistical Physics* 151(3), 623–636.
- Hillenbrand, A. and S. Hippel (2019). Strategic Inattention in Product Search. *MPI Collective Goods Preprint* (2017/21).
- Hofbauer, J. and K. Sigmund (2003). Evolutionary Game Dynamics. *Bulletin of the American Mathematical Society* 40(4), 479–519.
- Holmström, B. R. (1978). *On Incentives and Control in Organizations*. Stanford University.
- Holmström, B. R. (1982). *On the Theory of Delegation*. Northwestern University.
- Inderst, R. and M. Ottaviani (2009). Misselling through Agents. *American Economic Review* 99(3), 883–908.
- Inderst, R. and M. Ottaviani (2012a). Financial Advice. *Journal of Economic Literature* 50(2), 494–512.
- Inderst, R. and M. Ottaviani (2012b). How (not) to pay for advice: A framework for consumer financial protection. *Journal of Financial Economics* 105(2), 393–411.
- Jacquet, N. L. and S. Tan (2007). On the Segmentation of Markets. *Journal of Political Economy* 115(4), 639–664.
- Jagadeesan, M. and A. Wei (2018). Varying the Number of Signals in Matching Markets. In *International Conference on Web and Internet Economics*, pp. 232–245. Springer.
- Juul, J., A. Kianercy, S. Bernhardsson, and S. Pigolotti (2013). Replicator dynamics with turnover of players. *Physical Review E* 88(2), 022806.
- Kadam, S. V. (2015). Interviewing in Matching Markets. *Working Paper*.
- Lee, R. S. and M. Schwarz (2017). Interviewing in two-sided matching markets. *The RAND Journal of Economics* 48(3), 835–855.

- Li, J., N. Matouschek, and M. Powell (2017). Power Dynamics in Organizations. *American Economic Journal: Microeconomics* 9(1), 217–41.
- Lipnowski, E. and J. Ramos (2020). Repeated delegation. *Journal of Economic Theory* 188, 105040.
- Loertscher, S., E. V. Muir, and P. G. Taylor (2022). Optimal market thickness. *Journal of Economic Theory* 200, 105383.
- MacMillan, D. and N. Anderson (2019). Student tracking, secret scores: How college admissions offices rank prospects before they apply. *The Washington Post*. October 14. <https://www.washingtonpost.com/business/2019/10/14/colleges-quietly-rank-prospective-students-based-their-personal-data/> (accessed November 24, 2022).
- Mukhopadhyay, A. and S. Chakraborty (2020). Periodic Orbit can be Evolutionarily Stable: Case Study of Discrete Replicator Dynamics. *Journal of Theoretical Biology* 497, 110288.
- Nash, J. F. (1950). Non-cooperative games. Ph.D. diss., Princeton University. Reprinted in *The Essential John Nash*, ed. H. W. Kuhn and S. Nasar, 51-83, Princeton University Press, 2002.
- Palacios-Huerta, I. (2003). Professionals Play Minimax. *The Review of Economic Studies* 70(2), 395–415.
- Park, I.-U. (2005). Cheap-Talk Referrals of Differentiated Experts in Repeated Relationships. *RAND Journal of Economics* 36(2), 391–411.
- Romanyuk, G. (2017). Ignorance is Strength: Improving the Performance of Matching Markets by Limiting Information. *Harvard University Cambridge, Working Paper*.
- Roth, A. E. (2018). Marketplaces, Markets, and Market Design. *American Economic Review* 108(7), 1609–58.
- Roth, A. E. and X. Xing (1997). Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists. *Journal of Political Economy* 105(2), 284–329.

- Samuelson, L. (1988). Evolutionary Foundations of Solution Concepts for Finite, Two-Player, Normal-Form Games. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*, pp. 211–225.
- Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.
- Schottmüller, C. (2019). Too good to be truthful: Why competent advisers are fired. *Journal of Economic Theory* 181, 333–360.
- Shimer, R. and L. Smith (2001). Matching, Search, and Heterogeneity. *The B.E. Journal of Macroeconomics* 1(1).
- Sobel, J. (1985). A Theory of Credibility. *Review of Economic Studies* 52(4), 557–573.
- Sobel, J. (2013). Giving and Receiving Advice. *Advances in economics and econometrics* 1, 305–341.
- Sorin, S. (2020). Replicator dynamics: Old and new. *Journal of Dynamics & Games* 7(4), 365.
- Spiekermann, S., J. Grossklags, and B. Berendt (2001). E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pp. 38–47. ACM.
- Tsai, J. Y., S. Egelman, L. Cranor, and A. Acquisti (2011). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22(2), 254–268.
- Van Long, N. (2010). *A Survey of Dynamic Games in Economics*, Volume 1. World Scientific.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Walker, M. and J. Wooders (2001). Minimax Play at Wimbledon. *American Economic Review* 91(5), 1521–1538.
- Weibull, J. W. (1997). *Evolutionary Game Theory*. MIT Press.

Xu, Y., K. Wang, B. Zhang, and Z. Chen (2007). Privacy-enhancing personalized web search. In *Proceedings of the 16th international conference on World Wide Web*, pp. 591–600. ACM.

Zimmer, M. (2008). Privacy on Planet Google: Using the Theory of "Contextual Integrity" to Clarify the Privacy Threats of Google's Quest for the Perfect Search Engine. *J. Bus. & Tech. L.* 3, 109.