

Essays on applications of simple behavioural rules : an evolutionary approach

Citation for published version (APA):

Khan, A. (2014). *Essays on applications of simple behavioural rules : an evolutionary approach*. [Doctoral Thesis, Maastricht University]. Maastricht University. <https://doi.org/10.26481/dis.20140703ak>

Document status and date:

Published: 01/01/2014

DOI:

[10.26481/dis.20140703ak](https://doi.org/10.26481/dis.20140703ak)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Essays on Applications of Simple Behavioural Rules: An Evolutionary Approach

Dissertation

to obtain the degree of Doctor at
Maastricht University,
on the authority of the Rector Magnificus,
Prof. dr. L.L.G. Soete,
in accordance with the decision of the Board of Deans,
to be defended in public
on Tuesday, July 03, 2014, at 12:45 hours

by

Abhimanyu Khan

Supervisor:

Prof. dr. P.J.J. Herings

Co-supervisor:

Dr. R. Peeters

Assessment Committee:

Prof. dr. H.J.M. Peters (Chairman)

Prof. dr. W.H. Sandholm (University of Wisconsin)

Dr. R. Saran (Yale-NUS College)

Dr. E. Tsakas

Acknowledgement

The past four years at Maastricht University have been an immensely rewarding experience. I am eternally grateful to my advisors Jean-Jacques Herings and Ronald Peeters for their guidance and supervision at each and every moment throughout the past four years. I have been very fortunate to have advisors who have been very generous (to the extent of probably being guilty of indulgent) with the most valuable resource i.e. time. I have received a wealth of advice and assistance from them. I have benefitted from Hans Peters, William Sandholm and Rene Saran for their comments on my thesis. I am thankful to William Sandholm for hosting me for a semester at the University of Wisconsin-Madison and the Maastricht METEOR International travel grant for making this visit possible. I would like to express my gratitude to Matthew Embrey, Arkadi Predtetchinski, Arno Riedl, Kaj Thomsson and Elias Tsakas for their feedback and input, especially during the job-market season, and to Sylvia Beenens, Nicole Hulsman and Elke Lucas for all things sundry. I would like to thank Severine vanden Eynde, Christine Gutekunst, Anant Joshi, Zsombor Meder and Mukul Tyagi for life outside the university. Even while not being in Maastricht, these four years in Maastricht would not have been as they have been without the support of my parents, wife, sister and family. I remain enriched and forever indebted at the same time.

Contents

1	Introduction	1
2	Cognitive Hierarchies	7
2.1	Introduction	7
2.2	Preliminaries	11
2.3	The evolutionary Nash bargaining game	14
2.4	General class of bimatrix games	18
2.5	Extensions	20
2.5.1	Alternative specification of level- k reasoning	21
2.5.2	Knowledge of utilities	22
2.6	Concluding Discussion	26
3	A Duopoly Imitation model	29
3.1	Introduction	29
3.2	Model, notation and standard results	32
3.3	Results	34
3.3.1	Unperturbed learning dynamic and absorbing states	34
3.3.2	Perturbed learning dynamic and stochastic stability	35
3.4	Discussion	43
3.5	Conclusion	45
3.6	Appendix	46
4	A Multi-firm Imitation model	55
4.1	Introduction	55
4.2	Model	57
4.3	Results	60
4.3.1	Nash equilibrium	60
4.3.2	Imitation, evolution and stochastic stability	61
4.3.3	Qualitative features of the (long-run) equilibrium	65

4.4	Discussion	67
4.5	Conclusion	70
4.6	Appendix	71
5	Local Imitation and Coordination	75
5.1	Introduction	75
5.2	Model and Preliminaries	80
5.3	Results	84
5.4	Discussion	90
5.5	Conclusion	94
5.6	Appendix	95
	References	103
	Nederlands Samenvatting	109
	Short Curriculum Vitae	111

Chapter 1

Introduction

The broad research theme of the collection of essays in this thesis concerns the effect of the use of simple behavioural rules by economic agents in strategic situations, and it is an attempt to contribute to economic applications of evolutionary game theory. Evolutionary game theory explores the long run outcomes obtained when economic units use heuristics or simple behavioural rules to choose actions or strategies, receive payoffs based on chosen actions and after having received information about payoffs, revise their actions based on these heuristics. The use of such rules of thumb may be due to complexity of a situation, lack of complete information, or some form of bounded rationality and the focus is on the long-run outcome that emerges as individuals encounter the same strategic situation recurrently.

This approach is complementary to traditional economic analysis that is primarily based on rationality of economic agents. Rationality, as used in economics, carries the import that decision making units act “optimally”, even in situations that are strategic in nature or involve uncertainty. This implicitly invokes strong assumptions on information availability and the agents’ ability to process information optimally. One justification that is often provided is that while economic units might not behave “rationally” right from the start, non-optimal behaviour would get corrected for, and over time, these adjustments would lead to outcomes predicted by the rationality model. Thus, while an adjustment process over time is hinted at, it is often not modeled explicitly, and this is where an evolutionary approach complements the more traditional approach.

This evolutionary perspective may be enriched by incorporating the influence of networks on social and economic interactions. Networks, by virtue of determining the structure of interactions and how individuals receive information, may affect payoffs and the choice of strategies over time. This brings various interesting questions to the fore-front: how do various network architectures affect the resulting outcome? Is it possible to influence the network structure to nudge the population or society to a preferable outcome? My thesis

aims to examine some of the issues alluded to above, and in this thesis, I focus on two behavioural rules, namely, the best response dynamic and the imitation dynamic.

One may think of best-response dynamics as the most obvious manner of weakening the assumption of rational behaviour leading to utility or payoff maximisation in strategic situations. Indeed, starting at least from Brown (1951) and fictitious play, it has a long tradition of being studied. Fictitious play is conceptualised by having individual players play a best-response to the empirically observed entire history of past play. However, in some situations, it may be more realistic to assume that the medieval and recent history of play is more germane than ancient history of play for in the decision-making of individuals. The framework of adaptive play as introduced by Young (1993, 1998) proceeds along this line. Individuals have access to a finite recent history of play and gather information by sampling observations from this history, and then choose a best-response to the information thus acquired.

Interestingly, one may think of fictitious or adaptive play on one hand, and common knowledge or rationality on the other, as two ends of the spectrum. Best-responses arising from fictitious or adaptive play may be thought of as the minimal depth of reasoning in strategic situations while common knowledge as a model of infinite depth of reasoning about strategic situations. It is very well possible that cognitive layers, or finite depths of reasoning, exist between these two extremes. One example of such middling depth of reasoning is when an individual anticipates the adaptive-play generated best response of the co-player and plays a best response to that. Similarly, yet higher levels of reasoning may exist, and this serves as the motivation for the second chapter of the thesis.

In Chapter 2, I study the effect of weakening rationality and the associated notions of common knowledge and infinite depth of reasoning. I restrict attention to population games where the basic behavioural trait is described by a (myopic) best-response to actions and strategies used in the past, and build on this depths of reasoning based on this ‘primitive’ best-response behavioural trait. I argue that each population may have individuals with differing levels of cognition. For example, an individual may choose a higher order best response, i.e. in anticipation that a simple best response to past actions will be chosen, he may play a best response to the simple best response – this can be thought of as a second order best response. Individuals may play even higher order best responses by choosing a best response to a second order best response, and so on. The objective is to bridge the gap between infinite depth of reasoning at one end of the spectrum and myopic best-response to a sample of past play on the other extreme. This is done by comparing the long-run outcome in population games with cognitive hierarchies to the long-run outcome in population games where behaviour is described by myopic best-response to a sample of past play (as in Young’s (1993) adaptive play). I find that higher cognition may have

a bearing on long run outcomes, but only when the possibility of mistakes or occasional deviations from the said behaviour is incorporated, and counter-intuitively, the change in the long-run outcome may not be in favour of – and can actually go against – the population consisting of individuals with higher cognition.

The following chapters are an examination of another well-known and well-studies dynamic: the imitation dynamic, wherein individuals imitate the actions that they observe or perceive to be more successful. The imitation dynamic shares a close connection with replicator-type dynamics where strategies or actions that are more profitable than average grow in the population. This has a clear-cut motivation and interpretation in the biological literature – more successful traits tend to spread in a population. Interestingly, from an economics point of view, Schlag (1998) shows that imitation of successful practices is also the behavioural rule that is expected to yield the highest payoff in the class of all bounded rationality rules with finite memory. Under appropriate conditions, there appears to be experimental support for this behavioural trait. In my thesis, I assume that individuals imitate the strategy of the most successful individual, the use of which may be guided by the salience and prominence of successful individuals.

In Chapter 3 and Chapter 4, I study the behaviour of firms that imitate the strategies of those with the highest profit in the previous period. This is inspired by Alchian (1950) who argues that firms (owing to informational constraints) may only be relative profit-maximising, as this provides a ready reference point by which a firm can judge its own performance. A firm may consider imitating the most successful firm to avoid being left behind. While there is a branch of literature that explores this possibility, it is mainly restricted to firms in a homogeneous market. For example, Vega-Redondo (1997) and Alós-Ferrer and Ania (2000) study imitation by quantity-setting firms and price-setting firms respectively. The analysis of imitation driven firm behaviour in these chapters is in the context of differentiated markets. Further, adopting an evolutionary perspective allows us to attempt opening of a black-box in firm theory – that a firm may choose either price or quantity. I allow firms to choose price and quantity in advance, leaving open the possibility of non-market clearing. This stands in contrast to the standard Cournot model where firms choose quantities in advance but the determination of prices is somewhat a black box or the Bertrand model, where firms choose prices but produce on demand.

In Chapter 3, I analyse a duopoly situation where two firms are located at the extreme ends of a Hotelling line. It is possible that the long-run outcome of the imitation process (when firms imitate the more profitable firm but occasionally make mistakes) results in non-Nash equilibrium outcomes where the outcomes may be non-market clearing and prices may be higher than monopoly price. The intuition behind the observance of non-Nash equilibrium outcomes is that if a firm chooses an action that reduces its profit but leaves the rival even worse off, its actions are imitated. Such spiteful actions play a dominant role

in driving the long-run outcome under imitation. The potency of spiteful actions is also the reason for the emergence of the Walrasian outcome, which is characterised by a lower price than the Nash equilibrium, as the long run outcome in Vega-Redondo (1997). In contrast to Vega-Redondo (1997), I find that prices higher than the Nash equilibrium price may be the manifestation of spiteful actions. Noteworthy is the fact that the relevance of the role of spite is not limited to evolutionary economics; Hamilton (1976) serves to highlight the role of spite in an evolutionary context in the biology literature.

Interestingly, in Chapter 4, when the imitation model is extended to include more firms (on a Salop circle), the long-run outcome with imitation coincides with the Nash equilibrium. This stands in contrast to both the results of Chapter 3 and to a larger literature where the spite effect is built in an imitation dynamic through the focus on comparative performance. In this chapter, the spatial nature of competition results in competition amongst firms being local (i.e. primarily with the neighbouring firms), whereas imitation is global. This lends stability to the Nash equilibrium, even with an imitation dynamic. It would be interesting to extend this framework to incorporate heterogeneous behaviour – for example, firms may also be myopic best responders (i.e. they best respond to past behaviour). Heterogeneity can be conceived of in two different ways – each firm can choose to imitate or best respond, or each firm can either be only an imitator or only a best-responder with both types of behaviour existing in the population. It may be seen that the result obtained is robust the first type of heterogeneity mentioned above (i.e. with the possibility of within firm heterogeneity), and it may be instructive to examine the effect that the second type of heterogeneity (i.e. across firm heterogeneity) may have on the emergent outcome.

Chapter 5 examines the importance of information in a coordination game. The important issue in coordination games is that of efficiency – while coordinating on any action is an equilibrium, coordinating on a particular action may be better than coordinating on any other action. However, lack of information about the state of the world (read actions chosen by others) may be a confounding factor in this process of attaining coordinated actions. To model the feature that an individual may not be aware of the state of the world, I assume that each individual has a social network, and that each individual is aware of the state of the network. However, an individual may be randomly matched to interact with another individual in the population, who may not be a part of the former’s network. Thus while an individual receives information about the payoffs of different actions and strategies from others in his network, he may have to interact with an individual about whom he has no information. Imitation is restricted to the most successful individual in his neighbourhood. I analyse network structures that facilitate or hinder the process of coordination to the socially desirable outcome. I find that if each individual is sufficiently well connected to others in the population, then coordination on the most preferred action

occurs, but if even a few individuals are secluded or not connected well enough, then society may end up coordinating on an inferior outcome. The intuition for this result is that less desirable practises may originate and spread from segregated parts of the population, and that a minimum level of observability acts a bulwark against the invasion by less desirable practises. This hints at the benefit that more fluid information flows may confer on the entire population.

In my thesis, I have attempted to ask and provide a preliminary answer to some questions, but this process has led me to other intriguing questions, which may provide fodder for future research. There are two particular questions that I would like to make mention of. The first potential research question, similar in spirit to heterogeneously behaving firms alluded to above, deals with an expansion on the theme of heterogeneously behaved agents. Economic agents might perceive a given situation differently, or even if the perception is the same, they may be motivated to respond in systematically different manners. It may be instructive to incorporate this aspect to analyse economic phenomena and explore the effect of heterogenous behavior on the long-run; for example, how does an imitation driven system differ from a system where agents may imitate or best-respond? Further, do some behavioral types play a dominant role in giving direction to the long run outcome? For example, in a system where mixed behavioural rules co-exist, does one particular rule primarily decide the long-run resting point of the process? The difficulty of this approach is that one can conceive of various behavioural traits that may be relevant in a particular situation and numerous ways of incorporating such behavioural traits. The challenge lies in developing a reasonable and tractable method of introducing the above mentioned heterogeneity that may permit clear conclusions to be drawn.

The second potential research question lies at the interface of evolutionary game theory and networks, with the focus being the network characteristics that are critical in affecting the long-run outcome in strategic situations. An approach that may be particularly helpful in this endeavour is simulation analysis of the strategic situation under consideration. For instance, it may be instructive to use this approach to examine the outcome of coordination games on networks, where the role of the network might determine either the information structure (as in Chapter 5) or interaction structure. One may simulate such a network which, depending on the application in mind, could be formed randomly formed or through a process of more active deliberation by the nodes (i.e. players involved). The purpose is to explore if there are network characteristics that are pivotal in giving direction to long-run outcomes for a wide variety of behavioural rules and strategic situations. The advantage of this approach may be that it permits the study of complex strategic situations and it may be more amenable to introduction of heterogeneity in behavioural rules. Finally, one may hope that such simulation studies may not only complement theoretical analyses but may inspire new theoretical constructs.

Chapter 2

Cognitive Hierarchies¹

2.1 Introduction

One of the primary focuses of game theory is to model individual decision making in settings of strategic interdependence. The more traditional workhorse employed to this end has been the assumption of rationality and common belief of it, which at times might be a very demanding requirement. Binmore (1987) discusses fundamental issues with this assumption, including the inconsistency that it cannot be replicated by a Turing machine. Furthermore, experiments on human subjects reveal that for certain applications, the assumptions of common belief of rationality and its associated implication of an infinite depth of reasoning appear too strong, resulting in a disparity between theoretical predictions and empirical observations in the laboratory. This has generated considerable interest in a paradigm where individuals are either capable of, or make use of finite depth of reasoning. This could be either because of cognitive limitations on individuals or because of a belief that the co-player in the strategic situation will employ finite depth of reasoning, in which case it is also optimal to do so.

To illustrate this, consider the guessing game experiment of Nagel (1995) wherein individuals had to guess the number in the interval $[0, 100]$ which would be the closest to $p < 1$ times the mean of all guesses. Under common belief of rationality, the only rationalizable choice is to guess 0. The empirical data of the experiment, though, suggests that only a handful of the subjects guess 0. A considerable amount of the chosen strategies can be described by what is called “level- k ” behavior. Here, level-1 behavior refers to best responding to the belief that the mean of the guesses is equal to 50. This could be either due to the salience of the number 50 (Schelling salience) or because of the belief that the guesses of the others can be approximated by a uniform distribution over the interval

¹This chapter is based on the paper “Cognitive hierarchies in adaptive play”, Meteor Research Memorandum 12/007, and is co-authored with Ronald Peeters.

$[0, 100]$. In either case, the optimal guess arising out of level-1 behavior is $50p$. Level-2 behavior supposes that the others will employ level-1 reasoning and guess $50p$. Consequently, it is optimal to guess $50p^2$. Higher order level- k behavior is defined iteratively with level- k behavior being identified with guessing $50p^k$. Nagel (1995) finds that most of the subjects exhibit level-1 and level-2 behavior. The decisions in subsequent repetitions show a declining trend in the guesses. Two mutually non-exclusive explanations are an increase in the subjects' depth of reasoning and a shift in reference point for the most primary belief. Nagel's (1995) data does not provide any evidence for the former explanation and leans strongly towards the second explanation: subjects seem to use the mean of the previous period as the reference point.

A fair amount of theoretical and experimental studies has been directed towards the development of models that explicitly take into account the beliefs of individuals about other individuals' strategic decisions. For example, Stahl (1993) and Stahl and Wilson (1995) present a theoretical model of such players. The latter use a series of experiments to test for this pattern. Camerer et al. (2004) develop a cognitive hierarchy model in which each player assumes that she is the most sophisticated type and that the other individuals are of lower cognition. These models seem to explain the pattern of play in initial rounds of experiments – such as the one of Nagel (1995) – fairly well. They also are found to be robust to a wide specification of games: see, for instance, Stahl and Wilson (1994, 1995), Crawford and Iribarri (2007), Wang et al. (2010) and Coricelli and Nagel (2009).

Our interest, on the other hand, is to examine the effect of such a hierarchical model of players on the long-run process. Thus, not only are we interested in a framework where we begin with the hypothesis that individuals come in varying degrees of sophistication, but allow the different levels of cognition to operate on past experience as indicated by the empirical data of play. While the sophistication of a particular individual does not change over time, individuals pay attention to the way in which the game has unfolded in the recent past. The importance of modeling this feature is also emphasized in Binmore (1988). Although his main focus is on, what he calls, the “eductive process” (i.e., the cognitive process of the individual during the decision-making) rather than the “evolutive process”, he recognizes the importance of the latter as well. In this chapter, we consider a process which, while evolutionary in nature, also has the flavor of the eductive process. We believe that Young's (1993a, 1998) model of adaptive play provides a very convenient platform to incorporate these elements.

In Young's model of adaptive play, the interaction is modeled as a population game, where a population of players is associated with a specific role in a game. Populations have, at their disposal, a record of recent history of play. In each period, an individual is drawn randomly from the parent population to play the game. The individual samples

incompletely and without replacement from the past play of the rival population and plays a best response to the empirical distribution of strategies as revealed by the sample. In terms of the level- k model, this means that the individuals, by best responding to the empirical distribution available to them, exhibit level-1 behavior. Using an argument similar to Hurkens (1995), Young (1998) demonstrates that play converges almost surely to a minimal curb set.² The minimal curb sets form the recurrent classes (absorbing sets) of the dynamic process described by simple best response behavior. In order to select amongst the minimal curb sets (which may be numerous), the notion of stochastic stability is adopted.

We introduce the element of sophistication in this model of adaptive play in order to study the effect of higher cognition on long run outcomes. Young's model, with only simple (best responding) players, serves as the point of reference. We postulate that populations are comprised of individuals of varying levels of sophistication.³ The most primitive behavioral trait is described by the level-1 individuals referenced to above.⁴ A level-2 individual holds the belief that the rival is a level-1 type. Consequently, she samples from the actions of the own population to estimate the best response of a level-1 individual in the rival population and best responds to the estimate just formed. A level-3 individual plays a best response to the belief that the rival is of level-2. Hence, the level-3 individual samples from the actions of the rival population to estimate the best response of a level-1 individual in the own population. This is then used to estimate the best response of a level-2 individual in the rival population. The level-3 individual plays a best response to this estimate of the strategy that might be used by the level-2 individual of the rival population. This process is iterated one step more for each higher level of sophistication. So, as in Nagel (1995), each player assumes that the other players are one step down on the cognitive scale.

In Proposition 2.1, we show that the convergence to a minimal curb set, as shown by Young (1998) for the setting with only simple players, is invariant to the introduction of any higher level of sophistication in any population. A refinement of the minimal curb sets of the game may be obtained by considering an infinitesimally small probability of mistakes or experimentation, whereby individuals choose any strategy (and not necessarily a best response of some order). This selects the set of stochastically stable states. The possibility

²The minimal curb set is, loosely speaking, a set which contains all its best responses and there is no proper subset contained in it with the same property. For a more precise definition, see Section 2.2.

³Mohlin (2012) shows that it is possible for evolutionary learning processes to converge to a state where different cognitive types co-exist.

⁴For the purpose of nomenclature, we retain the associated terminology of the level- k model, even though we step outside the boundaries of it. Our focus is the long-run behavior, whereas the level- k model is meant for the purpose of explaining initial play.

of mistakes or experimentation now makes transition across minimal curb sets possible and the minimal curb sets that are most “stable” to such mistakes comprise the set of stochastically stable states. In Propositions 2.2–2.4, we find that the stochastically stable set may be sensitive to the composition of cognitive types in the populations. Proposition 2.2 shows this for the specific instance of the evolutionary Nash bargaining game, which is used as the leading example, while Propositions 2.3 and 2.4 do so for generic games.

In Young (1993b) it is shown that in the evolutionary Nash bargaining game with simple players only, the long-run outcome is the generalized Nash bargaining solution with the bargaining power being proportional to the populations’ own sample size. The intuition is simple: the smaller the sample size, the higher is the responsiveness to mistakes in the other population, leading to a smaller bargaining power.⁵ In Proposition 2.2 we show that the introduction of more sophisticated individuals in only one of the populations has no effect on the long-run outcome in case this population has the larger sample size, but that the standard Nash bargaining solution (that is, equal bargaining power) is obtained as long-run outcome in case this population has a smaller sample size.⁶ This finding is driven by the fact that level-2 individuals best respond to a best response to the sample of their own population. Via their beliefs, the level-2 individuals, when endowed with a smaller sample size, make the rival population appear more responsive to mistakes made by the own population. We show that there are no benefits of adding individuals of even higher level of cognition. Moreover, we show that if both populations host sophisticated individuals, then we obtain the standard Nash bargaining solution as the stochastically stable state.

In Proposition 2.3 we show that for generic games (where we adopt Young’s (1998) notion of genericity), the introduction of a fraction of sophisticated individuals in one population has an effect on the stochastically stable set only if the population with these individuals has a smaller sample size.⁷ The effect of higher cognition is to “decrease” the sample size of the rival population to that of the own population. In line with this, Proposition 2.4 says that if both populations host individuals with at least level-2 cognition, then irrespective of the relative sample size, the stochastically stable set is the same as that in Young’s adaptive play model with both populations having equal sample size.

⁵Although in the Nash bargaining game a higher sample size confers a benefit to the population, depending on the payoff structure it can be a bane as well – see Section 2.6 for more details.

⁶This effect has also been shown by Sáez-Martí and Weibull (1999) when allowing for “clever” individuals (that is, of level-2) in only one population.

⁷Matros (2003) shows this for the situation where only one population has level-2 individuals and both populations have the same sample size, in which case, the presence of the level-2 individuals does not make a difference to the stochastically stable outcome. Our more general result allows for even higher levels of cognition, populations to have unequal sample sizes, and both populations to host individuals of level-2 or higher.

Thus, if we believe both populations to be heterogeneous in the level of sophistication, then difference in sample size across populations no longer affects the stochastically stable outcome.

In the following section, we specify the model of adaptive play with higher cognitive individuals in more detail and state that this process of adaptive play converges to a minimal curb set. In Section 2.3 and Section 2.4, we study the impact of higher cognitive types on the long-run outcome for the evolutionary Nash bargaining game and general games respectively. In Section 2.5, we provide a discussion on two possible alternative specifications of our model. In our model we assume that an individual of a particular level believes she is playing someone of one level lower. In the first alternative specification, we show that our results are robust to the alternative specification where she believes her rival to be of any lower level – as is assumed for instance in Stahl (1993) and Camerer et al. (2004). In the second alternative specification, for a specific class of games, we relax the knowledge (individuals within) populations may have about cardinal preferences of (individuals within) the rival population. As a result, higher cognitive types use their own utility function while forming a belief about the rival’s play. Binmore (1988) mentions that such form of introspection might be a plausible way in which higher cognition operates. We find that play still converges to a minimal curb set (as ordinal preferences do not change), but the stochastically stable set may change (as the sensitivity to mistakes may change): in the evolutionary Nash bargaining game, the generalized Nash bargaining solution with bargaining power proportional to the rivals’ sample size and the 50–50 split may become feasible long-run outcomes. Finally, we close with a discussion of the implication of our results in Section 2.6. Amongst other things, we discuss the qualitative impact of higher cognition. The distribution of strategies in the long-run outcome might be affected, possibly resulting in payoff differences across individuals within and across populations. Further it is, perhaps counter-intuitively, shown that an entire population can be worse-off by the presence of more cognitive individuals.

2.2 Preliminaries

There are two finite populations, A and B , and each population is assigned a specific role in a specified two-player game that is played recurrently between randomly selected individuals from the two populations. Each individual has a positive probability of representing the population that she belongs to.⁸ The selected individual from population A (B) has to choose a strategy x_A (x_B) out of a finite set of strategies X_A (X_B). The pair of chosen strategies $(x_A, x_B) \in X_A \times X_B$ yields a payoff of $\pi_A(x_A, x_B)$ and $\pi_B(x_A, x_B)$ to

⁸We do not require the probability to be selected to be equal for each individual in a population.

the individual from population A and B respectively. We assume each population to be homogenous, i.e. all individuals in a population have the same utility function. The only heterogeneity we assume within a population is in the level of sophistication.

We assume individuals play adaptively. Both populations have access to the strategies chosen in the last m periods of play and individuals in population A (B) can draw a sample of proportion a (b) – i.e., of length am (bm) – without replacement. We assume each possible sample to have a positive probability of being drawn, but do not require all samples in the history to be equally likely to be drawn.

We identify the empirical distribution over chosen strategies with level-0 ($L0$) behavior. So, a level-1 ($L1$) individual plays a best response to her sample from the rival population's strategies.⁹ Next, a level-2 ($L2$) individual holds the belief that her rival is $L1$ and attempts to best respond to the strategy that the $L1$ individual of the rival population might choose. As a result, the $L2$ individual samples from her own population's past play, formulates the best response to the drawn sample¹⁰, adopts this as estimate of the strategy that the rival $L1$ individual might use, and plays a best response to it. The behavior of any individual of higher level of sophistication is described analogously.

It is evident from the above that an Lk individual, with k odd, draws a sample from the strategies of the rival population. In contrast, an Lk individual, with k even, draws a sample from the strategies of the own population. Following that, the sample is processed in accordance to the cognition of the individual.

At this point, we make the assumption that if a population contains an Lk individual with $k \geq 2$, then it also contains an $L(k - 1)$ individual.¹¹ So, we do not allow the entire population to be composed of, say, $L4$ individuals only. And, in case a population contains an $L4$ individual, then it also contains at least one $L3$, $L2$ and $L1$ individual. This permits us to describe the cognitive types present in a population by the most sophisticated individual. We use \hat{k}_A and \hat{k}_B to refer to the highest cognitive level in population A and population B respectively.

Given any arbitrary history on m periods of play, the adaptive process described above yields a Markov process on the state space $\Omega = (X_A \times X_B)^m$, where the states are the possible histories of length m . Let this (unperturbed) process, for sample proportions a and b and highest levels of sophistication \hat{k}_A and \hat{k}_B , be denoted by $P^{m,a,\hat{k}_A,b,\hat{k}_B}(0)$.

Let C_i be a nonempty subset of X_i ($i = A, B$). We denote the set of probability distributions over C_i by $\Delta(C_i)$. Moreover, by $BR_i(C_j)$, we denote the set of strategies in

⁹In cases of multiple best responses, we always assume each best response to be chosen with positive probability, not necessarily with equal chance.

¹⁰In order to do so, it is necessary that the $L2$ individual possesses knowledge of the utility function of the rival population. In Section 2.5 we are going to relax this assumption.

¹¹In Section 2.5 we replace this assumption with an alternative one.

X_i that consists of the best replies to any mixture in $\Delta(C_j)$ for the individuals in population i ($j \neq i$). Now, we can define the notion of (minimal) curb sets, which is due to Basu and Weibull (1991). The product set $C = C_A \times C_B$ is *closed under best replies* (or C is a *curb set*) if $BR_A(C_B) \times BR_B(C_A) \subseteq C$. Such a curb set is *minimal* if it does not properly contain a curb set.

Proposition 2.1. *If the history m is sufficiently large and the sampling is sufficiently incomplete (i.e., a and b are sufficiently small), then the minimal curb sets are the recurrent classes of the process of adaptive play $P^{m,a,\hat{k}_A,b,\hat{k}_B}(0)$.*

Proof. The proof for this is essentially the same as in Young (1998, Thm. 7) and Matros (2003, Thm. 1). Young (1998) shows that best response adaptive play converges to a minimal curb set in a finite number of steps when there are only L_1 individuals in either population. Since in our setting there is for any finite number of periods, a positive probability that only L_1 individuals are chosen during these periods, we know that our process converges to a minimal curb set. Matros (2003) shows that a fraction of L_2 individuals in one population does not have any disruptive effect on the convergence to a minimal curb set. The idea is that, by definition of the minimal curb set, higher order best responses are contained in it. Using the same arguments, we can show that the same holds for L_k individuals in either population. After play converges to a minimal curb set (by successively drawing L_1 individuals), responses of all higher cognitive types are contained in it such that play never leaves it. Furthermore, it is possible to transit from one state in the recurrent class to another in finite time by considering the event that only L_1 individuals are chosen. The result then follows from Young (1998). ■

The model of adaptive play so far has been built on the assumption that individuals are simple or sophisticated best responders. This results in the convergence to a minimal curb set. So, minimal curb sets represent the recurrent classes of the (unperturbed) Markov process. In order to allow for transits from one minimal curb set to another, we need to perturb the process. For the perturbed process, we assume that each individual has a probability ε to experiment or to commit a mistake; that is, they may choose any strategy – even those that are not a best response or a best response of a best response and so on, to any conceivable sample drawn from the history of past play. These experimentations induce the resulting perturbed process of adaptive play $P^{m,a,\hat{k}_A,b,\hat{k}_B}(\varepsilon)$ to be ergodic, such that for each $\varepsilon > 0$ a unique stationary distribution exists; that is, a unique solution $\mu(\varepsilon)$ to the equation $\mu \cdot P^{m,a,\hat{k}_A,b,\hat{k}_B}(\varepsilon) = \mu$. The *stochastically stable set* is defined to consist of precisely those states that receive a positive weight in the limiting stationary distribution $\mu^* = \lim_{\varepsilon \downarrow 0} \mu(\varepsilon)$.

The stochastically stable set is a subset of the set of recurrent classes (or, in the present case, the minimal curb sets) of the unperturbed adaptive process $P^{m,a,\hat{k}_A,b,\hat{k}_B}(0)$. Intuitively, this set consists of the minimal curb sets that are easier to reach and more difficult to transit from with (a series of) experimentations. We now make an important observation. We are interested in the minimum number of experimentations that thereafter make a transition from one minimal curb set to another possible via the best-response dynamic. The best responses of Lk individuals, with k odd (even), are only affected by a sufficient number of experimentations in the rival (own) population. For higher cognitive types, an Lk individual will have a best response outside the prevailing minimal curb set only if she believes that the $L(k - 1)$ individual in the rival population will do so. But this happens if she believes that the $L(k - 1)$ individual believes that an $L(k - 2)$ individual will do so. Ultimately, this can be traced to the belief that an $L1$ individual in some population chooses a strategy outside the recurrent class. If this does not happen, then none of the higher cognitive individuals in any population believe that the actions of the rival is going to change. The belief that an $L1$ individual in the rival (own) population does so, induces a change in behavior by Lk individuals, with k odd (even).

In the upcoming section we illustrate the effect of the presence of higher order best responders in Young's evolutionary Nash bargaining game. However before doing so, we comment on what might be a potentially important issue. One might question the behaviour of the Lk individuals by suggesting that while they believe the opponent population consists entirely of $L(k - 1)$ agents, they may be likely to observe apparently confounding data from their sample of past play. We, however, would like to argue that this may need not be a contradiction as long as the individuals are aware that everyone samples incompletely from the past history of play. Recognising this, individuals may still desire to use an appropriately sophisticated best response as that is the best they may be able to do given the information they have. In addition, because of the recurrent nature of the game where each individual has only a positive probability of being chosen to play, it is possible that an individual does not even observe the potentially confounding data.

2.3 The evolutionary Nash bargaining game

The Nash bargaining game is a two-player game where each of the two players demands a portion of some good (or, amount of money). If the total amount requested by the players is less than that available, both players get their request. If their total request is greater than that available, neither player gets their request. In this section we consider the Nash bargaining game in an evolutionary framework, played recurrently by randomly selected players from two finite populations.

Let A and B be two populations consisting of a finite number of individuals. In each period t , one individual is selected at random from each population to play the Nash bargaining game. Simultaneously and independently, each individual announces a demand from the feasible set of demands $D(\delta) = \{\delta, 2\delta, \dots, 1 - \delta\}$, say x_A^t and x_B^t . Each of them receives their respective demand if the two demands sum up to not more than the whole (that we normalize to one); i.e., if $x_A^t + x_B^t \leq 1$. Otherwise, both receive nothing. We assume that all individuals in population A have the same concave, strictly increasing, differentiable von Neumann-Morgenstern utility function $u : [0, 1] \rightarrow \mathcal{R}$ with $u(0) = 0$. Similarly, we assume all individuals in population B to have a utility function $v : [0, 1] \rightarrow \mathcal{R}$ with the same properties.

At period t , individuals α and β are chosen from populations A and B . Both individuals have access to a record of play of the last m periods: $\omega^t = ((x_A^{t-m}, x_B^{t-m}), \dots, (x_A^{t-1}, x_B^{t-1})) \in \Omega = (D(\delta) \times D(\delta))^m$. Individual α can draw a sample of size $a m$ of demand pairs (x_A^τ, x_B^τ) from the last m periods of play without replacement. Similarly, individual β draws a sample of size $b m$ of demand pairs (x_A^τ, x_B^τ) from the last m periods of play without replacement. We do not require each pair to be drawn with equal probability, but it is essential that any sample of the appropriate size can be drawn. Next, individual α (β) makes a demand x_A^t (x_B^t) that maximizes her expected payoff against the empirical distribution of demands as given by the sample drawn (given her depth of reasoning).

A state $\omega \in \Omega$ is a *convention* if it consists of some fixed division $(\bar{x}_A, \bar{x}_B) \in D(\delta) \times D(\delta)$ repeated m times in succession and in addition, $\bar{x}_A + \bar{x}_B = 1$. In Theorem 1 of Young (1993b) it is shown – for populations solely consisting of $L1$ individuals – that if at least one individual in each population samples at most half of the record of play, then (from any initial state) the process converges almost surely to a convention. This is actually a corollary of Proposition 2.1, as the set of conventions coincides with the set of all minimal curb sets. In the remainder of this section, we will focus on the (set of) stochastically stable state(s) as the precision of demand δ goes to zero.

In Young (1993b) it is shown that when both populations consist solely of $L1$ individuals ($\hat{k}_A = \hat{k}_B = 1$), the stochastically stable state is the generalized Nash bargaining solution with bargaining power proportional to the sample sizes. The generalized Nash bargaining solution is defined to be the convention given by the division (\bar{x}_A, \bar{x}_B) that maximizes the Nash product $(u(x_A))^a(v(x_B))^b$ subject to $x_A, x_B \in [0, 1]$ and $x_A + x_B = 1$. The standard Nash bargaining solution is simply the generalized Nash bargaining solution with $a = b$, or following Young's interpretation, it is the generalized Nash bargaining solution with both populations having equal bargaining power.

In the following proposition, we first present the result when both populations consist of only $L1$ individuals (as in Young (1993b)). Secondly, we allow the highest level of

cognition in population A to be larger than or equal to two: $\hat{k}_A \geq 2$. Thirdly, we allow the possibility of clever (and even more sophisticated) agents in the B population as well: $\hat{k}_B \geq 2$. Throughout, we maintain the assumption that the history of play, m , is sufficiently large and the sample (i.e. a, b) that individuals draw is sufficiently incomplete and examine the stochastically stable states as the precision of demand δ goes to zero.

We summarize all results in the following proposition, where (i) is due to Young (1993b) and a part of (ii) is covered by Sáez-Martí and Weibull (1999).

Proposition 2.2. (i) Suppose that both populations comprise of $L1$ individuals. The stochastically stable state is the generalized Nash bargaining solution with the bargaining power of each population equal to the sample proportion of each population.

(ii) Suppose population A has a fraction of sophisticated best responders. Then, the stochastically stable state is the generalized Nash bargaining solution if $a > b$. Otherwise, the stochastically stable state is the standard Nash bargaining solution.

(iii) Suppose both populations have sophisticated best responders. The stochastically stable state is the standard Nash bargaining solution.

Proof. With unperturbed adaptive play, the play converges to a convention. This is because the minimal curb sets of the game coincide with the conventions. To identify the stochastically stable state, we need to look at the minimum number of experimentations needed to transit from one convention to another.

Part (i). When there are only $L1$ individuals in each population, then the minimum number of mistakes needed to move away from the convention $(x, 1 - x)$ with positive probability corresponds to the smallest integer greater than or equal to $m r^\delta(x)$, where

$$r_A^\delta(x) = a \min\left\{\frac{u(x) - u(x-\delta)}{u(x)}, \frac{u(x)}{u(1-\delta)}\right\}, \quad (2.1)$$

$$r_B^\delta(x) = b \min\left\{\frac{v(1-x) - v(1-x-\delta)}{v(1-x)}, \frac{v(1-x)}{v(1-\delta)}\right\}, \text{ and} \quad (2.2)$$

$$r^\delta(x) = \min\{r_A^\delta(x), r_B^\delta(x)\}. \quad (2.3)$$

To explain this, note that there are two ways that we can move away from a convention. One is when individuals from population B make mistakes and these mistakes are sampled by individuals in population A . The other is when individuals from population A make mistakes and these mistakes are sampled by individuals in population B . Equation (2.1) captures the former possibility; Equation (2.2) the latter.

There are two mistakes an individual of population B can make: demanding an amount more than $1 - x$ and demanding an amount less than $1 - x$. Amongst the demands greater than $1 - x$, the least resistant one is to ask for δ more: i.e., demanding $1 - x + \delta$. A best responding individual of population A would shift to asking $x - \delta$, when this mistake has

been made sufficiently frequent. Let ℓ represent the least number of times population B has to make this mistake in order to displace the convention with positive probability. Given that population A individuals draw a sample of size $a m$, they will shift to asking $x - \delta$ when $\frac{am-\ell}{am} u(x) \leq u(x - \delta)$. This gives the least proportion of mistakes $\frac{\ell}{m} = a \frac{u(x)-u(x-\delta)}{u(x)}$, which is one of the terms in Equation (2.1). The second term comes from the situation when population B demands less than $1 - x$ by mistake. Among these mistakes, demanding δ is the least resistant one. Equation (2.2) can be derived in a similar manner.

Consider the fraction

$$\frac{r^\delta(x)}{\delta} = \min\left\{a \frac{u(x)-u(x-\delta)}{\delta u(x)}, a \frac{u(x)}{\delta u(1-\delta)}, b \frac{v(1-x)-v(1-x-\delta)}{\delta v(1-x)}, b \frac{v(1-x)}{\delta v(1-\delta)}\right\}. \quad (2.4)$$

As $\delta \downarrow 0$, the first term can be written as $a \frac{u'(x)}{u(x)}$ and the third as $b \frac{v'(1-x)}{v(1-x)}$, while the second and fourth terms become unbounded. So, $\lim_{\delta \downarrow 0} \frac{r^\delta(x)}{\delta} = \min\left\{a \frac{u'(x)}{u(x)}, b \frac{v'(1-x)}{v(1-x)}\right\}$. The convention with the highest minimum resistance is the convention $(\bar{x}, 1 - \bar{x})$ with \bar{x} maximizing the latter expression. As $\frac{u'(x)}{u(x)}$ is decreasing in x and $\frac{v'(1-x)}{v(1-x)}$ is increasing in x , the maximum is at the (unique) solution to $a \frac{u'(x)}{u(x)} = b \frac{v'(1-x)}{v(1-x)}$. This is precisely the first-order condition for the maximization of $(u(x))^a (v(1 - x))^b$; i.e., the generalized Nash bargaining solution.

Part (ii). When population A contains Lk individuals with $k \geq 2$, while population B only contains $L1$ individuals, then the minimum number of mistakes needed to move away from the convention $(x, 1 - x)$ corresponds to the smallest integer greater than or equal to $m r^\delta(x)$, where

$$r_A^\delta(x) = a \min\left\{\frac{u(x)-u(x-\delta)}{u(x)}, \frac{u(x)}{u(1-\delta)}\right\}, \quad (2.5)$$

$$r_B^\delta(x) = \min\{a, b\} \min\left\{\frac{v(1-x)-v(1-x-\delta)}{v(1-x)}, \frac{v(1-x)}{v(1-\delta)}\right\}, \text{ and} \quad (2.6)$$

$$r^\delta(x) = \min\{r_A^\delta(x), r_B^\delta(x)\}. \quad (2.7)$$

With respect to Part (i), only the second equation has changed. The reason is that Lk individuals draw their sample from their own population's past play if k is even, while they draw their sample from their rival population's past play if k is odd. Since $\hat{k}_A \geq 2$, both types of individuals are present in population A . Now it is not only the individuals from population B who may respond to mistakes by population A ; individuals from population A themselves may also do so. For instance, an $L2$ individual from population A takes a sample from their own population, estimates from this sample the possible choices by the population B individual, and subsequently responds to that. This opens up the possibility of the $L2$ individual responding to the mistakes made by individuals in the own population. In case their sample is smaller (i.e., $a < b$) they may estimate a shift in population B whereas none of the individuals in population B would already consider a shift from the convention.

Note that for $a > b$, Equation (2.6) is identical to Equation (2.2), and we obtain the same outcome as in Part (i). Next, consider the case $a \leq b$, for which

$$\frac{r^\delta(x)}{\delta} = \min\left\{a \frac{u(x)-u(x-\delta)}{\delta u(x)}, a \frac{u(x)}{\delta u(1-\delta)}, a \frac{v(1-x)-v(1-x-\delta)}{\delta v(1-x)}, a \frac{v(1-x)}{\delta v(1-\delta)}\right\}. \quad (2.8)$$

Now, $\lim_{\delta \downarrow 0} \frac{r^\delta(x)}{\delta}$ simplifies to $a \min\left\{\frac{u'(x)}{u(x)}, \frac{v'(1-x)}{v(1-x)}\right\}$, which is maximized at the (unique) solution to the equation $\frac{u'(x)}{u(x)} = \frac{v'(1-x)}{v(1-x)}$, that precisely characterizes the standard Nash bargaining solution.

Part (iii). Suppose we have sophisticated best responders in both populations. Now, there exists an individual in each population who samples the past actions of the rival population and there exists an individual in each population who samples the past action of its own population. The minimum number of mistakes needed to move away from the convention $(x, 1 - x)$ corresponds to the smallest integer greater than or equal to $m r^\delta(x)$, where

$$r_A^\delta(x) = \min\{a, b\} \min\left\{\frac{u(x)-u(x-\delta)}{u(x)}, \frac{u(x)}{u(1-\delta)}\right\}, \quad (2.9)$$

$$r_B^\delta(x) = \min\{a, b\} \min\left\{\frac{v(1-x)-v(1-x-\delta)}{v(1-x)}, \frac{v(1-x)}{v(1-\delta)}\right\}, \text{ and} \quad (2.10)$$

$$r^\delta(x) = \min\{r_A^\delta(x), r_B^\delta(x)\}. \quad (2.11)$$

Suppose, without loss of generality that $a \leq b$. Then,

$$\frac{r^\delta(x)}{\delta} = a \min\left\{\frac{u(x)-u(x-\delta)}{\delta u(x)}, \frac{u(x)}{\delta u(1-\delta)}, \frac{v(1-x)-v(1-x-\delta)}{\delta v(1-x)}, \frac{v(1-x)}{\delta v(1-\delta)}\right\}, \quad (2.12)$$

such that $\lim_{\delta \downarrow 0} \frac{r^\delta(x)}{\delta}$ simplifies to $a \min\left\{\frac{u'(x)}{u(x)}, \frac{v'(1-x)}{v(1-x)}\right\}$. From Part (ii) we know that this expression is maximized at the standard Nash bargaining solution. ■

Young (1993b) shows, in a two-population world with only L_1 individuals, that the bargaining power of a population is determined by its sample size, thereby implying that populations with smaller sample sizes are in a disadvantageous position. Parts (ii) and (iii) of Proposition 2.2 show that the presence of sophisticated individuals, levels out the disadvantage caused by drawing smaller samples. In particular, when both populations comprise of higher cognitive types (i.e., L_k individuals with $k \geq 2$), any disadvantage borne out of unequal sample size is “corrected” for and both populations end up with equal bargaining power. Hence, sample size is no longer a consideration in the long-run outcome.

2.4 General class of bimatrix games

In this section, we examine adaptive play between populations containing individuals of higher cognitive level in generic bimatrix games, where we adopt Young’s (1998) notion

of generosity. We use the stochastically stable set of Young's adaptive play model (i.e., with $L1$ individuals only) as a point of reference and denote this set by $\bar{\Omega}^{a,b}$ for population A and B with sample sizes a and b respectively. As before, we assume that the memory (m) is sufficiently large and the samples drawn are sufficiently incomplete, i.e. a and b are sufficiently small relative to m .

The next proposition deals with the case where only one of the populations (population A) comprises of higher cognitive types.

Proposition 2.3. *Suppose that the highest level of cognition in population A is Lk with $k \geq 2$, while population B only contains $L1$ individuals. The stochastically stable set is described by $\bar{\Omega}^{a,b}$ if $a \geq b$ and by $\bar{\Omega}^{a,a}$ otherwise.*

Proof. From Proposition 2.1, we know that curb sets are invariant to cognitive hierarchies and that the unperturbed adaptive play process converges to a curb set.

Let ℓ be the minimum number of experimentations or mistakes required to transit from one recurrent class to another when there are only $L1$ individuals in either population. Moreover, let ℓ' be the minimum number of experimentations or mistakes required to transit from one recurrent class to another when there are only $L1$ individuals in population B while population A contains more sophisticated individuals (i.e., $\hat{k}_A \geq 2$).

The set of stochastically stable states might change if the minimum number of experimentations needed to transit from one recurrent class to another changes. Notice that if it is possible to get from one recurrent class to another in ℓ experimentations when $\hat{k}_A = 1$ and $\hat{k}_B = 1$, then it is certainly possible as well with ℓ experimentations when $\hat{k}_A \geq 2$ and $\hat{k}_B = 1$, since with positive probability only $L1$ individuals are chosen in population A . Therefore $\ell' \leq \ell$. The converse holds as well, i.e. $\ell \leq \ell'$. To see this, note that Lk individuals, with $k \geq 2$ odd, would play a strategy outside the prevailing minimal curb set only if there are enough number of mistakes in the rival population's past play. Thus, the minimum number of mistakes required to induce a best response outside the prevailing minimal curb set does not change. As a result, there are no transitions that take place in the presence of Lk individuals, with $k \geq 2$ odd, that would not take place with $L1$ individuals. Lk individuals, with $k \geq 2$ even, sample from their own population's past play and it would take a certain number of mistakes in their own population's past strategies for these individuals to play a best response not contained in the prevailing minimal curb set.

When sample sizes are equal (i.e., $a = b$), the same number of mistakes are sufficient to induce $L1$ individuals in population B to play a different strategy. Hence, the same transition can be established in a process with only $L1$ individuals. When population A has a larger sample size (i.e., $a > b$), $L1$ individuals in population B would in fact,

owing to the smaller sample size, respond to a lesser number of mistakes in population A . The minimum number of experimentations for the transition involve $L1$ individuals only. Thus, the possibility of a similar transition also exists with only $L1$ individuals in either population. So, the set of stochastically stable states is $\bar{\Omega}^{a,b}$ if $a \geq b$.

Now, suppose population A has a smaller sample size (i.e., $a < b$). Then, Lk individuals in population A , with $k \geq 2$ even, react to the mistakes that appear in their own population earlier than the individuals in population B . The $L1$ individuals in population B would have reacted to the same mistakes if they had a sample size equal to a . As a result, the same transition would be effected with only $L1$ individuals in both populations, where both populations are endowed with sample size a . Hence, the set of stochastically stable states is $\bar{\Omega}^{a,a}$ if $a < b$. ■

Proposition 2.3 implies that the presence of individuals with higher cognition only influences the long run outcome when the population that they belong to is endowed with a smaller sample size. In such a case, the effect of these individuals with higher cognition is to “decrease” the sample size of the rival population to that of the own population.¹² As a consequence, when both populations comprise of higher cognitive types, the net effect is to “equalize” the sample size of both populations. This is precisely what the following proposition states. Since it is an implication of the previous proposition, we do not provide a formal proof.

Proposition 2.4. *Suppose both populations have individuals with a level of cognition of at least 2. The stochastically stable set is $\bar{\Omega}^{a,a} = \bar{\Omega}^{b,b}$.*

Thus the above two propositions generalize the result obtained for the evolutionary Nash bargaining model.

2.5 Extensions

In this section we present two extensions of the model considered. The first considers an alternative specification of level- k reasoning, where we assume an individual of level k to believe the opponent may be of any lower level with positive probability. In the second we assume that individuals of higher cognition (falsely) use their own utility function when forming a belief on the rival population’s play.

¹²The situation where one population contains a share of $L2$ individuals and both populations have an equal sample size has been captured in Matros (2003).

2.5.1 Alternative specification of level- k reasoning

In this first extension, we replace the assumption that Lk individuals believe that the opponent is of level $k - 1$ and that each population comprises a positive mass of any lower type than the maximum. Instead, we assume now that an Lk individual believes that the opponent is of at most level $k - 1$, with positive mass on all cognitive levels up to $k - 1$.¹³ Let us denote this individual by L^*k . Note that the Lk individuals defined earlier have point-beliefs; i.e. they believe that the rival is of type $L(k - 1)$ and place probability one on this belief. Their beliefs do not recognize that the rival's cognitive type could be, for instance, $L(k - 2)$. In contrast, the L^*k individual acknowledges the fact that the rival could be of any level between and including 1 and $k - 1$ with positive probability. In doing so, such individuals admit greater uncertainty in the distribution of their belief about the opponent and consider all best responses to all possible opponents, i.e. she constructs a best response to a $L(k - 1)$ individual, a best response to a $L(k - 2)$ individual and so on. In the end, one such best response is chosen and we assume that any such constructed best response has a positive probability of being chosen. Here, it is important to stress on the fact that while the L^*k individual considers the best-response to all lower cognitive types, at the point of choosing an action, the L^*k individual places point belief on one such lower cognitive type.¹⁴

This alternative specification of level- k reasoning does not affect the results stated in Propositions 2.3 and 2.4 of the preceding section. To see this, it suffices to realize that since an L^*k individual places a positive weight on all lower cognitive types, she behaves like an Lk (when a best response corresponding to the belief that the opponent is an $L(k - 1)$ individual is chosen) or an $L(k - 1)$ individual (when a best response corresponding to the belief that the opponent is an $L(k - 2)$ individual is chosen), each with positive probability. Hence, the minimum number of experimentations to transit from one minimal curb set to another does not change in comparison to the model with Lk individuals.

Recall that we do not require L^*k individuals to assign precise probabilities on the likelihood to meet a rival of particular (lower) level of cognition, as is done in Stahl's (1993) model of sophistication and in Camerer et al.'s (2004) model of cognitive hierarchies. The latter assumes individuals to perfectly know the fraction of lower cognitive types in the population, such that the beliefs of individuals with higher level of cognition are closer to the actual proportions in the population. It is straightforward to see that these restrictions

¹³Note that we do not explicitly require all these types actually to be contained in the rival population.

¹⁴The L^*k individual, therefore, does not choose a best response to a distribution of types, but rather, after considering the best response to each lower cognitive type, places point mass belief on one such lower cognitive type. We remain agnostic about the process by which the particular lower cognitive type is chosen, but only require that the probability of each lower cognitive type being assigned point mass belief be strictly positive.

on the beliefs, do not affect the recurrent classes (which are the minimal curb sets) and the stochastically stable set as long as the perceived probability to meet a type of any lower cognitive level is strictly positive.

2.5.2 Knowledge of utilities

One source of dissatisfaction with the model in the previous sections might be that individuals with level of cognition higher than or equal to 2 know the utility function of the individuals in the other population. This knowledge allows more cognitive individuals to estimate the probable best-responses of the rival. In this subsection, we relax this assumption for the class of games G' , where the strategy sets and the ordinal preferences are identical for all individuals in either population.¹⁵ We assume that individuals only know their own (population's) utility function and not that of their possible rivals. So, even though the ordinal preferences are the same across the two populations, individuals in one population are not aware of the cardinal utilities of the individuals in the other population. This implies that if for population A with utility function u , $u(a, b) > u(c, d)$ for any pure strategies a, b, c and d in the common strategy set, then for population B with utility function v , $v(a, b) > v(c, d)$ holds, and vice-versa. We denote the resulting individual with level- k depth of reasoning by $L'k$. Note, though, that we retain the assumption that if a population contains an $L'k$ individual ($k \geq 2$), then it also contains an $L'(k-1)$ individual. Let \hat{k}'_A and \hat{k}'_B denote the highest cognitive level in population A and population B respectively.

Like Lk individuals, $L'k$ individuals draw samples from the strategies of the rival population when k is odd and of the own population when k is even. So, $L'1$ individuals behave identically to $L1$ individuals: they best respond to the empirical distribution of strategies in the sample drawn from the opponent population's past play. $L'2$ individuals best respond to their estimate of their rival's behavior. Assuming that their rival is an $L'1$ individual, they sample from their own population's past play to assess the strategies the rival may be using. However, for this assessment they use their own utility function rather than that of the rival.¹⁶

The following proposition claims that the resulting unperturbed adaptive play process is invariant to this modification of the level- k model.

Proposition 2.5. *In the class of games G' , if the history m is sufficiently large and the sampling is sufficiently incomplete (i.e., a and b are sufficiently small), then the process of adaptive play converges to a minimal curb set with probability one.*

¹⁵ Examples outside this class for which the propositions below do not hold are easily constructed.

¹⁶ Possible explanations of such a systematic behavioral trait include the ‘false consensus effect’ (Ross et al., 1977) and self-projection (Buckner and Carroll, 2007).

Proof. The proof is a trivial modification of the proof of Proposition 2.1. With positive probability, only $L'1$ individuals are chosen from either population till play converges to a minimal curb set. We now only have to show that play cannot leave this minimal curb set, or in other words, for any $L'k$ individual from either population, the best-reply is contained in the support of the set of strategies that forms the minimal curb set. We show this specifically for the best-response used by $L'2$ individuals. The reasoning for other $L'k$ individuals is similar.

Let $BR_A(C) \subseteq S$ be a *minimal* set for population A that contains *all* pure strategy best responses to all (pure and mixed) strategies in set $C \subseteq S$. If population A individuals sample from any history of set C and play a best-response to the sample, then the best-response will be contained in $BR_A(C)$. On the other hand, when $L'2$ individual from population B samples from set C and uses his own utility function (i.e. v) to estimate the best response of population A, then we will show that this estimated best response for the other population is also in $BR_A(C)$. (This ensures that the $L'2$ individual will choose a strategy that is a best-response to some strategy in $BR_A(C)$, and so play does not leave the minimal curb set.) But for now, let the latter set be labeled $BR_B(C)$, i.e. let $BR_B(C) \subseteq S$ be a *minimal* set for population A that contains all pure strategy best responses to all (pure and mixed) strategies in set $C \subseteq S$. This implies that for any (mixed) strategy $q \in C$, $\{\text{argmax}_r u(r, q), \forall r \in S\} = BR_A(C)$. Let the support of q above be (q_1, \dots, q_n) and the associated probabilities be (p_1, \dots, p_n) . Then,

$$p_1 u(r, q_1) + \dots + p_n u(r, q_n) > p_1 u(t, q_1) + \dots + p_n u(t, q_n), \forall r \in BR_A(C), \forall t \in S \setminus BR_A(C)$$

or,

$$p_1(u(r, q_1) - u(t, q_1)) + \dots + p_n(u(r, q_n) - u(t, q_n)) > 0, \forall r \in BR_A(C), \forall t \in S \setminus BR_A(C)$$

By the identical ordinal preferences assumption, $\text{sgn}(u(r, q_l) - u(t, q_l)) = \text{sgn}(v(r, q_l) - v(t, q_l))$, for all $l = 1, \dots, n$ (where sgn is the sign function). By continuity of real numbers, we can say that there exists a probability distribution (p'_1, \dots, p'_n) over (q_1, \dots, q_n) , such that,

$$p'_1(v(r, q_1) - v(t, q_1)) + \dots + p'_n(v(r, q_n) - v(t, q_n)) > 0, \forall r \in BR_A(C), \forall t \in S \setminus BR_A(C)$$

or,

$$p'_1 v(r, q_1) + \dots + p'_n v(r, q_n) > p'_1 v(t, q_1) + \dots + p'_n v(t, q_n), \forall r \in BR_A(C), \forall t \in S \setminus BR_A(C)$$

Hence, it must be that $r \in BR_B(C)$. As this holds for any $r \in BR_A(C)$, we can infer that $BR_A(C) \subseteq BR_B(C)$. But now we can use the exact same reasoning to say that if $w \in$

$BR_B(C)$, then $w \in BR_A(C)$, and so, $BR_B(C) \subseteq BR_A(C)$. Thus, $BR_A(C) = BR_B(C)$. Thus, the best-response chosen by $L'2$ individuals (of either population) will be contained in the minimal curb set. Further, it can be easily verified that this also holds true for populations that contain more sophisticated individuals. ■

For further selection among the recurrent classes (i.e., the minimal curb sets) of the unperturbed process, we again adopt the notion of stochastic stability by studying the support set of the limit of the stationary distribution of regular perturbations of the process when the perturbation vanishes. Despite the invariance of the recurrent classes with respect to the current alternative specification of level- k behavior, the stochastically stable set may change. The next proposition illustrates how the stochastically stable state of the evolutionary Nash bargaining game is changed.

Proposition 2.6. (i) Suppose that population A contains a fraction of $L'k$ individuals for $k \geq 2$, while population B is comprised by $L'1$ individuals only. The stochastically stable state is either the generalized Nash bargaining solution with the bargaining power of a population equal to its own sample proportion, or the 50–50 division (equal split).

(ii) Suppose that both populations comprise of a fraction of $L'2$ individuals. The stochastically stable state is either the generalized Nash bargaining solution with the bargaining power of a population equal to its own sample proportion, a modified generalized Nash bargaining solution with the bargaining power of a population equal to the other population's sample proportion, or the 50–50 division.

Proof. By Proposition 2.5, play settles in a minimal curb set, which corresponds to a convention. For stochastic stability we need to examine the relative ease or difficulty of transiting from a convention.

Part (i). A convention $(x, 1 - x)$ can be disrupted by mistakes by population A or by mistakes by population B .

First, consider mistakes by population A . One possibility is that these mistakes affect the sample of $L'1$ individuals in population B . For these individuals to change their demand the minimum number of mistakes required equals $r_{BA}^\delta(x) = b \min\{\frac{v(1-x)-v(1-x-\delta)}{v(1-x)}, \frac{v(1-x)}{v(1-\delta)}\}$. The other possibility is that these mistakes are sampled by $L'k$ individuals in population A with k even. For these individuals to change their demand the minimum number of mistakes required equals $r_{AA}^\delta(x) = a \min\{\frac{u(1-x)-u(1-x-\delta)}{u(1-x)}, \frac{u(1-x)}{u(1-\delta)}\}$. After all, these individuals would demand something apart from x if they believe that an $L'(k-1)$ individual in population B is going to demand something else than $1 - x$. The latter possibility occurs if the $L'(k-1)$ individual in population B believes that an $L'(k-2)$ individual in population A is going to demand something different than x , and so on. Ultimately, this can be traced to the belief that an $L'1$ individual in population B changes the demand. If this does not happen, then

none of the higher cognitive individuals in any population believe that the action of the rival is going to change. Individuals in population A form this belief when they sample a sufficient number of mistakes in their own population. Since we are looking at the minimum number of mistakes to move out of a convention, it suffices to examine the case when the mistakes in their own population are such that the mistaken demands are just more than x (i.e., $x + \delta$) or when they ask for the least amount (i.e., δ). Suppose that ℓ mistakes of the first type have been made. For an individual A to consider it viable that population B may change (to $1 - x - \delta$), ℓ should be such that $u(1 - x - \delta) \geq \frac{am - \ell}{am} u(1 - x)$. This gives the least proportion of mistakes necessary to be $\frac{\ell}{m} = a \frac{u(1-x) - u(1-x-\delta)}{u(1-x)}$. The second term can be derived by considering the second type of mistakes.

Next, consider mistakes by population B . These mistakes can only affect the sample of $L'k$ individuals in population A with k odd. For these individuals to change their demand the minimum number of mistakes is required to be $r_{AB}^\delta(x) = a \min\{\frac{u(x) - u(x-\delta)}{u(x)}, \frac{u(x)}{u(1-\delta)}\}$.

The minimum resistance of the convention $(x, 1-x)$ is thus given by $r^\delta(x) = \min\{r_{BA}^\delta(x), r_{AA}^\delta(x), r_{AB}^\delta(x)\}$. It is easily seen (by using the same argument as in the proof of Proposition 2.2) that $\lim_{\delta \downarrow 0} \frac{r_\delta(x)}{\delta} = \min\{a \frac{u'(x)}{u(x)}, a \frac{u'(1-x)}{u(1-x)}, b \frac{v'(1-x)}{v(1-x)}\}$. The first of these terms is decreasing in x while the latter two are increasing in x . The stochastically stable efficient allocation $(\bar{x}, 1 - \bar{x})$ is at the (unique) \bar{x} that maximizes $\frac{r_\delta(x)}{\delta}$, which is at the unique solution to $a \frac{u'(x)}{u(x)} = \min\{a \frac{u'(1-x)}{u(1-x)}, b \frac{v'(1-x)}{v(1-x)}\}$. Depending on which of the two terms in the right hand-side is lowest at \bar{x} , we either find the 50–50 division, or the generalized Nash bargaining solution with the power of the population equal to its own sample size.

Part (ii). A convention $(x, 1 - x)$ can be disrupted by mistakes by population A or by mistakes by population B .

First, consider mistakes by population A . One possibility is that these mistakes affect the sample of $L'k$ individuals in population B with k odd. For these individuals to change their demand the minimum number of mistakes required equals $r_{BA}^\delta(x) = b \min\{\frac{v(1-x) - v(1-x-\delta)}{v(1-x)}, \frac{v(1-x)}{v(1-\delta)}\}$. The other possibility is that these mistakes are sampled by $L'k$ individuals in population A with k even. For these individuals to change their demand the minimum number of mistakes required equals $r_{AA}^\delta(x) = a \min\{\frac{u(1-x) - u(1-x-\delta)}{u(1-x)}, \frac{u(1-x)}{u(1-\delta)}\}$.

Next, consider mistakes by population B . One possibility is that these mistakes affect the sample of $L'k$ individuals in population A with k odd. For these individuals to change their demand the minimum number of mistakes required equals $r_{AB}^\delta(x) = a \min\{\frac{u(x) - u(x-\delta)}{u(x)}, \frac{u(x)}{u(1-\delta)}\}$. The other possibility is that these mistakes are sampled by $L'k$ individuals in population B with k even. For these individuals to change their demand the minimum number of mistakes required equals $r_{BB}^\delta(x) = b \min\{\frac{v(x) - v(x-\delta)}{v(x)}, \frac{v(x)}{v(1-\delta)}\}$.

The minimum resistance of the convention $(x, 1-x)$ is thus given by $r^\delta(x) = \min\{r_{BA}^\delta(x), r_{AA}^\delta(x), r_{AB}^\delta(x), r_{BB}^\delta(x)\}$. It is easily seen (by using the same argument as in the proof of

Proposition 2.2) that $\lim_{\delta \downarrow 0} \frac{r_\delta(x)}{\delta} = \min\{a \frac{u'(x)}{u(x)}, a \frac{u'(1-x)}{u(1-x)}, b \frac{v'(x)}{v(x)}, b \frac{v'(1-x)}{v(1-x)}\}$. The first and third of these terms are decreasing in x while the other two are increasing in x . The stochastically stable efficient allocation $(\bar{x}, 1 - \bar{x})$ is at the (unique) \bar{x} that maximizes $\frac{r_\delta(x)}{\delta}$, which is at the unique solution to $\min\{a \frac{u'(x)}{u(x)}, b \frac{v'(x)}{v(x)}\} = \min\{a \frac{u'(1-x)}{u(1-x)}, b \frac{v'(1-x)}{v(1-x)}\}$. Depending on which of the two terms in the two sides of the equation are lowest at \bar{x} , we either find the 50–50 division, the generalized Nash bargaining solution with the power of the population equal to its own sample size, or the generalized Nash bargaining solution with the power of the population equal to the other population’s sample size. ■

Proposition 2.6 shows that this alternative specification of level- k behavior has a severe qualitative impact on the long run outcome.¹⁷ The $L'k$ specification allows for the equal split and an inverted version of the generalized Nash bargaining solution to be feasible long run outcomes. This latter is not a long run outcome in the Lk specification, while the former was found only when the cardinal utilities of the populations were identical up to an affine transformation.

2.6 Concluding Discussion

In this chapter we investigated the effect of sophisticated agents on adaptive play. Apart from the more usual way of defining sophisticated agents as higher order best responders (Lk individuals), we introduced an alternative notion that does not assume (higher order best responding) individuals to be aware of the utility function of the individuals in the rival population ($L'k$ individuals). We find that the unperturbed adaptive play process, with either notion of sophistication, converges to a minimal curb set. The intuition for this result is simple. With only $L1 = L'1$ individuals, best response adaptive play converges to a minimal curb set in a finite number of steps. So, since with positive probability only simple best responders are chosen for any finite number of steps, the process converges to a minimal curb set. By the construction and definition of a minimal curb set, higher order best responses are contained in it. Hence, any strategy chosen by a more sophisticated best responder would be in the minimal curb set.

In addition to the focus on the recurrent classes, we also considered the effect of these more sophisticated best responders on the stochastically stable outcomes relative

¹⁷Even though we assume that an $L'k$ individual projects his own utility function onto the rival population, the general message of this section – that is, play converges to a minimal curb set but the set of stochastically stable states may differ – are valid even if an $L'k$ individual evaluates the rival’s preferences with some other cardinal utility function (under the proviso that the ordinal preferences are identical). The reason for more explicitly dealing with self-projection of utility function is that under a situation of identical ordinal preferences, it might be more reasonable to attribute one’s own preference onto another rather than to use some other arbitrary utility function to do so.

to Young's setting with only simple best responders. We find that a higher level of sophistication has no effect when the population containing these has the same or higher sample size. However, sophistication has a differential effect when the population holding these individuals has a strictly smaller sample size. The effect is as if to reduce the sample size of the rival population to that of the own population. This effect is already obtained by the presence of a fraction of $L2$ individuals and the addition of even higher cognitive types has no further effect. Again, the intuition is simple. More sophistication has an influence only when it results in a formation of a belief that the rival is going to change her behavior, and the sophisticated best responder actually changes her own behavior on basis of that belief. An individual needs to be at least of level 2 to form such a belief. Only when the sophisticated best responders have a smaller sample size the belief is "unfounded", which is why that is the only situation in which they have a differential effect on stochastically stable outcomes. When both populations have $L2$ individuals, it is as if we are dealing with the basic adaptive play model of Young with both populations having equal sample size. The presence of $L2$ individuals on both sides, level out any differences in sample size. And, again, more sophistication beyond $L2$ types has no additional effect – as mentioned earlier, all higher beliefs formed can be traced to the behavior of an $L1$ type individual. With $L1$ and $L2$ individuals, the mistakes in either populations' play that may give rise to different beliefs are accounted for. Hence, even higher cognition has no effect.

We hazard a guess and suspect that these properties are transferable to a setting with multiple populations. First, convergence to a minimal curb set is clearly obtained for any number of populations and any distribution of cognitive types across populations. Moreover, regarding stability against experimentation, still all plays of strategies outside a prevailing curb set can be traced to either a level-1 individual in some population to change play with positive probability (based on a sufficient number of mistakes), or there should be a belief by some higher order type that this is about to happen.

Having identified the effect that higher cognition has on adaptive play, the next question is their qualitative impact: Do more sophisticated individuals leave their population better off? Thereto, suppose that population A consist of $L1$ and $L2$ individuals, while population B has $L1$ individuals only. Suppose further that the sample size of population B is greater than that of population A ($a < b$), so that by Proposition 2.4, the stochastically stable states are the same as in a process where both populations have simple best responders only and the sample sizes are equal. We argue that depending on the payoffs of the game, this outcome may be better or worse for population A – the population with the clever agents. That population A may benefit from the presence of clever agents was seen in the Nash bargaining model. However, this is not a universal phenomenon. To this end, we draw attention to an example from Young (1998, Ch. 5) that illustrates that having less

information (i.e., a smaller sample size and therefore higher sensitivity to mistakes) may be an advantage. Consider the game in Figure 2.1 and assume that the row population has a sample size of s while the column population has a sample size of s' , with $\frac{9}{19}s < \frac{7}{17}s'$. When both populations have simple best responders only, then the outcome with payoffs

	<i>L</i>	<i>R</i>
<i>U</i>	10, 7	0, 0
<i>D</i>	0, 0	9, 10

Figure 2.1: Example

(10, 7) is stochastically stable – i.e., the row population is better off. However, if the row population contains a fraction of sophisticated best responders (i.e., $L2$ individuals), then the outcome with payoffs (9, 10) is stochastically stable. Hence, the row population suffers from the presence of higher cognitive individuals in their population.

Another interesting qualitative impact of sophistication concerns the distribution of strategies inside the stochastically stable minimal curb set(s). In the game of Matching Pennies where one of the populations had a fraction of $L2$ individuals, Matros (2003) shows that in the long run outcome, the $L2$ individuals have a positive expected payoff while the $L1$ individuals in both populations have a negative expected payoff. Even though sophistication beyond $L2$ has no effect on the stochastically stable minimal curb set(s), we suspect that similar features may be observed, apart from a difference in the speed of convergence. A related question is, do more sophisticated types always end up with a payoff at least as large as that of less sophisticated types? We suspect that this will not generally be the case in a model where there are at least $L2$ types on both populations as the expected payoffs will depend heavily on the precise proportion of each type in each population.

Chapter 3

A Duopoly Imitation model¹

3.1 Introduction

The traditional literature in the field of oligopoly theory typically models firms as perfectly rational agents whose decisions are aimed at the maximization of profit. However, the recognition that agents might be informationally constrained and/or at most boundedly rational, has activated a growing interest in analyzing the impact that such constraints might have on decision making and market outcomes. Indeed, it is not unnatural to conceive of a situation where firms lack the relevant information for profit maximization. Hence, the tools usually used to study firm behavior loose their applicability and alternative approaches need to be adopted.

In this context, Alchian (1950) emphasizes the importance of relative performance and that relative profit might serve as the chief driver of firm behavior. If the salience of relative profit is accepted, then imitation of a more successful firm may be a reasonable behavioral assumption. This approach agrees with the evolutionary perspective, which is built on the premise that strategies which yield a higher-than-average payoff will tend to be more attractive. Huck et al. (1999), Huck et al. (2000), Offerman et al. (2002) and Apesteguia et al. (2007) find some experimental evidence of firms imitating the fitter firm once feedback on strategies and profits is explicitly provided.

In this spirit, Vega-Redondo (1997) studies an evolutionary Cournot situation where firms producing a homogeneous commodity compete in quantities and the market price is determined by the total quantity produced. The profits and quantities chosen by each firm are observable, and the action chosen by the most successful firm is almost always imitated. Occasionally, firms experiment and choose any quantity. As the probability of experimentation goes to zero, the Walrasian outcome prevails most of the time. Huck et al. (1999) and Apesteguia et al. (2010) provide experimental support for this finding.

¹This chapter is based on the paper “Evolution of Behaviour when Duopolists choose Price and Quantity”, Meteor Research Memorandum 11/027, and is co-authored with Ronald Peeters.

Similarly, Alós-Ferrer et al. (2000) analyze the Bertrand model in an evolutionary setting. Firms, with increasing marginal costs in a homogeneous market, set prices and the lowest price setters win the market. Here, analogous to Vega-Redondo (1997), prices announced by the firm and the profits realized by each are observed and the behavior of the firms is based on imitation of the successful firms. The study obtains a strict subset of the Nash equilibria of the underlying game as the long-run outcome of this dynamic process of imitation and occasional experimentation.

This chapter explores a related issue. The question that we seek to address is: What is the nature of a process where firms do not choose either price or quantity but both of them simultaneously? This is inspired by the accepted “shortcoming” of the Cournot and Bertrand model. In the former, firms choose quantity and market price is determined by market clearing while in the latter, firms choose price and production is on-demand. What if production occurs in advance (a la Cournot) and firms have the independence to choose their price (a la Bertrand)? It then seems reasonable to assume that firms choose both price and quantity.² This question is more interesting if we abandon the assumption of market clearing and then examine if it arises out of the evolutionary process.³

We consider a market for a good that is horizontally differentiated. To keep things simple, we assume that there are two firms that are located at the extreme ends of a Hotelling line.⁴ We impose a very basic informational setting; the only information that the firms have pertain to its own marginal cost and that there is a competing firm. In particular, we do not require the firms to be aware of the fact that they are in a differentiated market. Neither do they have to be cognizant of the consumers’ valuation of their product, the consumers’ preferences between the two firms (i.e. the transportation cost) and the size of the market.

The firms can choose both price and quantity. Market clearing is not imposed and hence the presence of excess supply or excess demand cannot be ruled out a priori. After setting prices and quantities, the firms realize profits, all three of which are observable. The behavior that we impose on the firms is that the less successful firm immediately imitates the more successful firm in both price and quantity. In addition to this imitation dynamic, we introduce the possibility of experimentation, whereby firms choose a price or quantity that is not necessarily dictated by imitation. The possibility of such experimentation

²For a slightly more detailed discussion on this, see, for example, Van den Berg and Bos (2012).

³Herings (1997) and more recently Alós-Ferrer and Kirchsteiger (2010) show that it is possible for non-market clearing institutions to evolve.

⁴The purpose of the assumption on firm location is simply to specify the demand that each firm faces. Hohenkemp and Wambach (2010) show that in an evolutionary context where a good has multiple attributes, firms would choose to locate at the centre of the Hotelling line. In our chapter, location is not a subject of choice and is given exogeneously to determine demand. Nonetheless, we discuss the implications of our results for homogeneous product markets in Section 3.4.

makes the system ergodic and enables us to obtain the long-run outcome. The techniques used to identify the long-run outcome are based on the concept of stochastic stability, as developed by Young (1993), Kandori et al. (1993) and Ellison (2000). This process of imitation and occasional experimentation shares parallels with Hamilton's theory of spite in biology (Hamilton, 1970). Via imitation, the system gravitates towards the action that yields higher relative profit, and this might occur at the expense of absolute profit.

Our results indicate that the outcome depends on the magnitude of the consumer's gross utility or valuation for the product. When the consumer valuation is lower than a certain bound, the unique stochastically stable state involves each firm behaving as a monopolist (Proposition 3.2). As consumer valuation increases up to another bound, we obtain a set of prices and quantities as the long-run outcome with all prices being above the monopoly price (Proposition 3.3). For sufficiently high consumer valuation, any state where firms choose identical price and quantity and receive at least zero profits can be a long-run outcome (Proposition 3.4). In this last case, in contrast to the ones mentioned earlier, market clearing does not necessarily come about – it is possible to have long-run outcomes with persistent excess supply or excess demand.

At this juncture, we provide some preliminary intuition for these results. When consumer valuation is sufficiently low, each firm is for all practical purposes, endowed with some monopoly power. This enables monopoly pricing to be more resistant to any spiteful behavior by the competing firm. But, with higher consumer valuations, this monopoly power is eroded. For an intermediate range of consumer valuation, the set of stochastically stable states are the ones which are most resistant to experimentation by a deviant firm. Strikingly, this goes along with prices above monopoly level. This is a reflection of the fact that the stochastically stable states are the ones most immune to any type of potential spiteful behavior by the competing firm – the high prices allow a firm to have a secure but admittedly smaller market. For sufficiently high consumer valuation, no state is more resistant to experimentation and all absorbing states are stochastically stable. Interestingly, in spite of all the emphasis on relative profit, excess supply (which represents unnecessary cost) or excess demand (which represents an opportunity to increase profit) are possible long-run features.⁵

The chapter is organized as follows. In Section 3.2, we present the model and the standard benchmarks. In Section 3.3, we present the main results. In Section 3.4, we discuss robustness of our results and comment on possible extensions, such as one wherein the quantity produced by the firm is not observable but the actual sales are. Section 3.5 concludes.

⁵The nature of the results here resonate with the principal message in Bergin and Bernhardt (2004), where it is reasoned that in an imitative dynamic, the decision-makers through their focus on comparative performance end up ignoring the externality effect of own actions. This often leads to payoffs being more depressed than they would be in a Nash equilibrium.

3.2 Model, notation and standard results

We study a duopolistic market in the standard Hotelling framework, but where firms choose prices and quantities. Two firms are located at the endpoints of the unit interval $[0, 1]$ along which the consumers are uniformly distributed. We identify firms and consumers by their location on this interval – firm 0 is the firm situated at location 0 and firm 1 is the firm situated at location 1. Two firms simultaneously and independently decide on the quantity to put on the market and the price to charge for each unit. After having observed the two prices, consumers choose from which of the two firms to acquire the product or to abstain from buying. Consumers acquire one unit of the good at most. All consumers assign a common intrinsic value of β to the good and on purchase, incur a transportation cost of τ per unit of distance to the chosen firm. Hence, the utility of the consumer at location $x \in [0, 1]$ equals

$$U_x = \begin{cases} \beta - \tau \cdot x - p_0 & \text{if } x \text{ purchases from firm 0;} \\ \beta - \tau \cdot (1 - x) - p_1 & \text{if } x \text{ purchases from firm 1;} \\ 0 & \text{if } x \text{ does not consider buying any good at all.} \end{cases}$$

Consumers choose the option that maximizes their utility. We assume that the price charged by a firm is more salient than quantity it produces and so consumers do not take into account the quantity produced by the firms. It is possible that the demand at one firm exceeds the supply while the other firm is left with an excess supply. We ignore the possibility of spillover demand – i.e. the possibility that a consumer who is not able to acquire the good at the preferred location travels to the other firm (which might have excess supply). Such considerations would require additional assumptions on how goods are rationed or distributed amongst consumers by the firm that faces an excess demand. To elaborate briefly, it is possible that a consumer prefers acquiring a good at either firm to not acquiring one at all while there might be another consumer who would prefer to acquire the good only from a specific firm. The possibility of spillover demand arises only if both these consumers prefer the same quantity-constrained firm and the second consumer gets served. To analyse this, there needs to be an assumption on which consumer acquires the good in case their preferred firm faces excess demand. While it may be instructive to look at this aspect, we do not pursue it here.

A firm's sale, s_i , is given by the minimum of the supply q_i and the demand d_i , where d_i is given by the mass of consumers that decides to purchase from firm i ($i = 0, 1$). We assume an equal and constant production cost of c per unit and that the goods are perishable such that excess supply can not be utilized. As firms are aware of their production costs, we assume they do not adopt prices below unit cost. A firm's profit is then given by:

$$\pi_i = p_i \cdot s_i - c \cdot q_i.$$

The specification of the demand of firm i depends on whether the prices give rise to all consumers deciding to purchase a good (irrespective of their success to acquire the good) or whether there is a positive mass of consumers that prefer not to purchase at given prices. In the former case the demand of the firms is determined by the location of the consumer that is indifferent between the two products: $\hat{x} = \frac{p_1 - p_0 + \tau}{2\tau}$. If $\hat{x} \geq 1$, then $d_0 = 1$ and $d_1 = 0$; if $\hat{x} \leq 0$, then $d_0 = 0$ and $d_1 = 1$; otherwise (that is, if $\hat{x} \in (0, 1)$) $d_0 = \hat{x}$ and $d_1 = 1 - \hat{x}$. On the other hand, when there is a positive mass of consumers that prefer not to purchase, the demands are given by $d_0 = \hat{x}_0$ and $d_1 = 1 - \hat{x}_1$, provided that these values are not below 0 or above 1. Here, \hat{x}_0 stands for the consumer who gets precisely zero utility if she purchases from firm 0, i.e. $\hat{x}_0 = \frac{\beta - p_0}{\tau}$; \hat{x}_1 is defined similarly, i.e. $\hat{x}_1 = 1 - \frac{\beta - p_1}{\tau}$. Note that in the former case $\hat{x}_1 \leq \hat{x} \leq \hat{x}_0$ and the firms actively compete for the indifferent/marginal consumer; in the latter case $\hat{x}_0 < \hat{x} < \hat{x}_1$ and the firms act as “local monopolists”. Also note that two firms compete for the marginal consumer (i.e. all consumers would receive a non-negative utility if they were able to purchase the good) if and only if, without loss of generality, $p_1 \leq 2\beta - p_0 - \tau$; this implies that if both firms choose the same price (say, p), then they compete for the marginal consumer if and only if $p \leq \beta - \frac{\tau}{2}$.

Throughout this chapter, we assume that the consumer valuation exceeds the unit production cost: $\beta > c$. This ensures that there is scope for efficient trade. Moreover, in the main body of the chapter, we focus on the situation where transportation costs are moderate relative to the marginal cost of production: $\tau < c$. This latter assumption is purely cosmetic and adopted to improve readability. However, the appendix contains the relevant results for the parameter configuration $\tau \geq c$.

Although we adopt an evolutionary approach where firms are not required to have any knowledge of the demand curves and the size of the market (or, alternatively, the location of the rival firm), for the purpose of benchmarking we provide some standard solutions that assume firms to possess the relevant information to maximize profits. Firstly, in case a firm is a monopolist located at one end of the Hotelling line, a price of $p = \max\{\frac{\beta+c}{2}, \beta - \tau\}$ would be chosen with a quantity equal to $q = \min\{\frac{\beta-c}{2\tau}, 1\}$. Secondly, the Nash equilibrium predicts the firms to produce this monopoly quantity and to charge the corresponding monopoly price when $\beta < c + \tau$. For $\beta \in [c + \tau, c + \frac{3}{2}\tau]$, Nash equilibrium prices and quantities chosen by each firm equal $p = \beta - \frac{\tau}{2}$ and $q = \frac{1}{2}$ respectively, while for $\beta > c + \frac{3}{2}\tau$, Nash equilibrium prices and quantities chosen by each firm equal $p = c + \tau$ and $q = \frac{1}{2}$ respectively.

3.3 Results

Instead of assuming firms to maximize profits, we postulate that firms learn by imitation and experimentation. The less fit firm (i.e. the firm with the lower profit) imitates the behavior of the fitter firm (i.e. the one with the higher profit). This requires strategic decisions and the yielded profits to be perfectly observable. Once firms settle on identical strategies, alternative price–quantity pairs may be chosen only by way of experimentation.

Note that for this learning dynamic, firms do not need all profit relevant information of the environment they compete in. In our setting, awareness of the size of the market, the intrinsic value consumers attach to a good and the transportation costs they are facing is not required. We even do not require the firms to be aware of their or their rival’s position on the Hotelling line. The only information we assume them to possess is their unit cost of production and the existence of a rival firm whom they can observe.

3.3.1 Unperturbed learning dynamic and absorbing states

We consider the following imitation dynamic. If in a particular state, a firm receives a negative profit, the rival firm is imitated in case the latter obtained a nonnegative profit. Otherwise, both firms shut down by choosing a production level of zero next period (without changing their price). If a firm receives a nonnegative profit, it imitates the rival firm in case this rival firm is strictly fitter in the current state (received a strictly higher profit); in case the rival firm is just as fit, we assume the firm imitates the rival with a positive probability strictly less than one. The state of the process is described by a vector (p_0, q_0, p_1, q_1) , where p_i and q_i is the price and quantity chosen by firm i . A state is said to be absorbing when there is no possibility of transiting to another state by imitation. The specified imitation dynamic gives rise to the following proposition.

Proposition 3.1. *The set of absorbing states are the monomorphic states that yield non-negative profit to the firms.*

Proof. A monomorphic state (i.e. a state where both firms choose the same price and quantity) with non-negative profit is absorbing as firms receive equal profits and have the same actions – this leaves no further scope for imitation. On the other hand, if we are in a state where firms realize unequal profits, then either both of them shutdown (this occurs when both firms have negative profit) or one firm imitates the other. In either case we transit to a monomorphic state. If we are in a state where firms realize equal profit with disparate actions, the positive probability of imitation makes this state transient. ■

Without experimentation the process will converge to an absorbing state. Experimentations that give the experimenting firm the higher profit relative to its rival are needed

to drive the system out of an absorbing state. The resistance of the absorbing state towards experimentation provides information on the stability of an absorbing state. In the next subsection we will characterize the set of states that are most resistant towards experimentation.

3.3.2 Perturbed learning dynamic and stochastic stability

Before specifying the experimentation probabilities, we begin by acknowledging a conceptual issue with experimentation in multiple dimensions, which is that there is no “theory” on how experimentations happen or occur. Papers such as Van Damme and Weibull (2002) and Blume (2003) tackle this concern by endogenizing the mistake probability by assuming that agents would be less likely to experiment away from states with good realizations. Such an approach is unfortunately not applicable here as firms in our model do not know the demand function that they face and, hence, cannot evaluate the payoffs corresponding to other choices.

Given this background, there are two features that (we feel) might be plausible to incorporate in the specification of experimentation. Firstly, since there are two firms who make decisions independently, it might be reasonable to speculate that experimentation by one firm is more likely than experimentation by two firms. Secondly, since firms act in two dimensions, it might be hypothesized that an experimentation in one dimension is more likely than an experimentation in two dimensions. Combining these two features, we assume that the probability of experimentation by one firm in either price or quantity (say, of the order of ε) is more likely than the experimentation by a firm in price and quantity (say, of the order of ε^2). The latter in turn is more likely than two firms simultaneously experimenting in price or quantity (say, of the order of ε^3).⁶ Due to experimentation, the perturbed process does not have any absorbing states; so, in the remainder of the chapter, when we mention absorbing states, we mean the absorbing states of the unperturbed process.

The subset of absorbing states that are most resistant to these experimentations possibilities are stochastically stable. Young (1993, 1998) and Kandori, Mailath and Rob (1993) develop a technique for identifying stochastically stable states that involves constructing, for each absorbing state, a tree such that there is a unique path from each of the other absorbing states. The resistance of an edge of the tree equals the minimum number of experimentations required for the transition along that edge (in our case, this number refers

⁶The results that we obtain depend partially on this assumption. This is not surprising in the light of Bergin and Lipman (1996), who show that the stochastically stable states are sensitive to the specification of how experimentations occur. Since there is no guiding principle on how experimentations occur, we adopt the system of experimentation as induced by the two features mentioned.

to the order of ε above). The sum of resistances of all edges gives the resistance of the tree. It is in general, possible to construct more than one tree for each absorbing state and so, the focus is, for each absorbing state, on the tree with minimum resistance. An absorbing state for which the minimum resistance tree has the maximum resistance is stochastically stable.

Based on the above procedure, Ellison (2000) develops an alternative method which involves radius, co-radius and modified co-radius of a set of absorbing states. The radius of a set of absorbing states is given by the minimum number of experimentations such that the unperturbed dynamic does not lead the process back to the initial set of absorbing states. For the co-radius of a set of absorbing states, we consider for each of the absorbing states outside the set, the minimum number of experimentations which, in addition to the unperturbed dynamic, are required to transit into into the set, from the absorbing state outside the set ; the maximum of the minimum resistance from other absorbing states outside the set is the co-radius of the set. Finally, for the modified co-radius of a set of absorbing states, we look (for each of the absorbing states outside the set) at the resistance of transition from another absorbing state; but now, if the path involved transition through other intermediate absorbing states, then the resistance from all these intermediate absorbing states is subtracted to obtain the “effective” resistance from an absorbing state to the set. The maximum of the “effective” resistance from all other absorbing states is the modified co-radius of the set. Ellison (2000) shows that if the radius of a set of absorbing states is greater than either the co-radius or the modified co-radius of the set, then the stochastically stable set is contained in that set. Further, for any two absorbing states, if the radius of the first absorbing state equals to minimum “effective” resistance of transition from the second absorbing state, and if the second absorbing state is stochastically stable, then so is the first absorbing state.

The set of stochastically stable states depends on the value of β , the gross utility that consumers receive from consumption. We present three propositions, each of them devoted to a distinctive case with respect to the magnitude of β .

Proposition 3.2. *Suppose $\beta < \tau - 2c + \sqrt{\tau^2 + 9c^2 - 2c\tau}$. The state where both firms supply the monopoly quantity $q^* = \frac{\beta-c}{2\tau}$ and set the monopoly price $p^* = \frac{\beta+c}{2}$ is the unique stochastically stable state.*

Proof. Let us denote the state mentioned in the proposition by ω^* . Since the supposition of the proposition implies $\beta < c + \tau$, the quantity produced by each firm according to ω^* is less than half the market size. Our proof is constructed as follows. Firstly, we show that no experimentation in price *and* quantity can push the system outside the basin of attraction of state ω^* . We then infer that it is also not possible with *either* one price

or one quantity experimentation. From this we conclude that the radius of ω^* is strictly larger than 2. Secondly, we show that the co-radius of ω^* is at most 2. From Theorem 1 in Ellison (2000), it follows that all stochastically stable states are included in the set $\{\omega^*\}$. Finally, existence leads us to conclude that ω^* is the unique stochastically stable state.

Part 1 (radius). In state ω^* the firms set price and quantity at the monopoly level and do not compete for the marginal agent. Experimentation in price by one firm leads to active competition with its rival if and only if the experimented price is less than $\frac{3}{2}\beta - \frac{1}{2}c - \tau$. Notice that any larger price leads to a decrease in the experimenting firm's profit, while it does not affect its rival's profit and will result, by imitation, in a return to ω^* . Suppose, therefore that the experimenting firm experiments with the price $p' < \frac{3}{2}\beta - \frac{1}{2}c - \tau$. Moreover, suppose that this firm experiments simultaneously with the optimal quantity given that price: $q' = \frac{\beta+c+2\tau-2p'}{4\tau}$. This experimentation leads to a profit of $\pi' = (p' - c)q'$ for the experimenting firm and a profit of $\pi = \frac{\beta+c}{2}(1-q') - c q^*$ for its rival. The difference between these two profits, $\pi' - \pi$, is maximized at $p'_{max} = \frac{c+\tau}{2}$. As $\tau < c$, this price is not feasible since it is below unit cost and so, it suffices to consider the case where the experimenting firm experiments with marginal cost pricing.

Suppose the experimenting firm experiments to $p' = c$ and $q' = \frac{\beta-c+2\tau}{4\tau}$. Then, $\pi' - \pi = \frac{1}{8\tau}(\beta^2 + 4c\beta - 2\tau\beta - 2c\tau - 5c^2)$, which is negative if and only if $\beta \in (\tau - 2c - \sqrt{\tau^2 - 2c\tau + 9c^2}, \tau - 2c + \sqrt{\tau^2 - 2c\tau + 9c^2})$. Since $\tau - 2c - \sqrt{\tau^2 - 2c\tau + 9c^2} < c < \beta$, and as $\beta > c$ by assumption, we get to the condition $\beta < \tau - 2c + \sqrt{\tau^2 - 2c\tau + 9c^2}$.

We have shown that even in the best case experimentation with prices and quantities, the experimenting firm cannot obtain a situation where its profit is equal or larger than that of its rival when the condition in the proposition is met. Now, note that if it is not possible to leave ω^* with a price experimentation in combination with an optimal quantity adjustment, it is also not possible without such an adjustment. On the other hand, an only quantity experimentation always decreases own profit (either because more units can be sold at a price above cost, or because more units are produced than are sold), while it doesn't affect the profit of the rival firm.

Part 2 (co-radius). From the first part of Part 1, we also learn that starting from any state, if a firm experiments simultaneously to price p^* and quantity q^* , then it results in a larger profit relative to the profit of its rival. This means that, from any state, at most two experimentations are needed to get to (the basin of attraction of) ω^* . ■

Proposition 3.2 reveals that when the intrinsic value of the product (β) is small, the monopoly outcome is expected to prevail in the long run. In particular, the two firms are not actively competing for the same customers and so, the market is only partially served. Notwithstanding the market not being fully served, the stochastically stable outcome im-

plies market clearing. Note, however, that the restriction on the intrinsic value is so severe that the participation constraint of the consumers at the other end of the market is not satisfied for a price equal to unit cost. This hands each of the firms some market power and in this framework of imitation and experimentation, it manifests itself into monopoly prices and monopoly outputs. It is interesting to note that both firms obtain the highest possible profit for this parameter configuration.

Proposition 3.3. *Suppose $\tau - 2c + \sqrt{\tau^2 + 9c^2 - 2c\tau} \leq \beta < c + \tau$. Then, the set of stochastically stable states consists of the monomorphic states with prices $p^* \in (\bar{p}, \beta)$ and quantity $q^* = \frac{\beta-p^*}{\tau}$, where \bar{p} is the unique feasible solution to $p(\frac{c-p+\tau}{2\tau}) - c\frac{\beta-p}{\tau} = 0$.*

Proof. Let Ω^* denote the set of stochastically stable states as it is claimed in the proposition. Notice that all states in Ω^* imply market clearing and in none of them firms actively compete for the marginal consumer as $\bar{p} \geq \frac{\beta+c}{2} > \beta - \frac{\tau}{2}$. Our proof is constructed as follows. Firstly, we show that the radius of Ω^* is 2. Secondly, we show that the modified co-radius of Ω^* is 1. It follows from Theorem 2 in Ellison (2000) that the stochastically stable states are contained in Ω^* . Finally, we show that all states in Ω^* are stochastically stable.

Part 1 (radius). We demonstrate that the radius of the set Ω^* is 2. We first show that the radius is at least 2 by showing that there exists no only-quantity experimentation or only-price experimentation that makes the system leave the basin of attraction of Ω^* .

Since supply equals demand in all states in Ω^* and the consumers' acquisition decision is independent of the quantity produced, a quantity experimentation reduces the own profit (as the price is above marginal cost) while it leaves the rival's profit unaffected. Hence, a quantity experimentation leaves the experimenting firm worse off as compared to its rival and imitation will drive the process back to the original state. So, there does not exist any disrupting only-quantity experimentation.

Let us now take up the case for price experimentation with a higher price. Since firms do not compete for the marginal agent, experimentation to a higher price ($p' > p$) will not lead to a change in the rival's profit. As the quantity is unadjusted, the cost component of the own profit does not change, and the revenue component is at least as much if $p'\frac{\beta-p'}{\tau} \geq p\frac{\beta-p}{\tau}$ – this is satisfied if and only if $p' \in [\beta - p, p]$; a contradiction.⁷ What about experimentation with a lower price? Experimentation with a lower price ($p' < p$) decreases the own profit, and can therefore only lead to a different state if the rival's firm is decreased by even more. For this to be the case, the price p' must be sufficiently low to compete for the rival's consumers: $p' \leq 2\beta - p - \tau$. Such a experimentation results in a situation where the experimenting firm obtains equal or larger profit than its rival if and

⁷Note that $\beta - p < p$, since $p > \bar{p} \geq \frac{\beta+c}{2} > \frac{\beta}{2}$.

only if it leads to equal or larger revenue (because costs are equal as the experimentation is only in price): $p'\frac{\beta-p}{\tau} - \frac{p'-p+\tau}{2\tau} \geq 0$. Since the derivative of the left-hand side to p' is the constant $(2\beta - 3p)$, it suffices to check the validity of the inequality at the two extreme values: c and $2\beta - p - \tau$. Experimentation to $2\beta - p - \tau$ can not be disruptive for the obvious reason that no consumers are stolen away and the own profit decreases. For $p' = c$ the inequality simplifies to $p^2 - (3c + \tau)p + 2\beta c \geq 0$. The derivative of the left-hand side equals $2p - (3c + \tau)$ and is negative since $p < \beta < c + \tau < \frac{3c + \tau}{2}$.⁸ This means that if the inequality is not satisfied for a certain price p , then it is not satisfied for larger values of p . The price \bar{p} is precisely the price at which the inequality is binding. Hence, for prices larger than \bar{p} , an experimentation to marginal cost is not disruptive.

We can conclude that for the states in Ω^* , the radius is at least 2. But now, notice that all states in Ω^* are with prices above monopoly level. An experimentation by one firm to monopoly price *and* quantity gives the experimenting firm a strictly higher profit and leads to an immediate drift towards an absorbing state outside Ω^* . Thus, the radius of Ω^* is 2.

Part 2 (modified co-radius). We now show that the modified co-radius of Ω^* is 1. We do so by demonstrating that from any other state outside Ω^* , it is possible to reach any state in Ω^* with a sequence of single suitable experimentations.

Firstly, let us suppose that the initial state is of the form $\omega^0 = (p, q, p, q)$ with $p < \beta - \frac{\tau}{2}$ such that firms compete for the marginal consumer. And, let $\omega^* = (p^*, q^*, p^*, q^*)$ be the state in Ω^* that we try to reach. If both firms earn non-negative profit state, then ω^0 is absorbing; otherwise firms stop producing and the process transits to state $(p, 0, p, 0)$. Irrespective of the sign of the profit, if the rival firm experiments to $q = \frac{1}{2}$, it will receive at least the same profit. This means that one experimentation is needed to get to the state $\omega^1 = (p, \frac{1}{2}, p, \frac{1}{2})$, which is an absorbing state as each firm makes non-negative profit. Next, consider an experimentation to marginal cost which would imply a larger profit relative to that of the rival firm if and only if the rival firm receives a negative profit: $p\frac{c-p+\tau}{2\tau} - \frac{c}{2} \leq 0$, which simplifies to $(p - \tau)(p - c) \geq 0$. This inequality will always be fulfilled as $p \geq c > \tau$, and the other firms (with positive probability) imitates marginal cost pricing. Hence, from ω^1 , one experimentation to marginal cost pricing is needed to get to the absorbing state $\omega^2 = (c, \frac{1}{2}, c, \frac{1}{2})$. Now, from ω^2 , consider an experimentation in quantity to q^* . Note that $q^* < \frac{1}{2}$, since $p^* > \bar{p} > \beta - \frac{\tau}{2}$. Both firms earn exactly zero profit when one firms experiments with q^* ; with positive probability quantity it gets imitated and we are in the absorbing state $\omega^3 = (c, q^*, c, q^*)$. Finally, from ω^3 , let one of the firms experiment with p^* . Then, this experimentation gets imitated with positive probability if

⁸The first inequality follows from p being in (\bar{p}, β) , the second by assumption of the proposition, and the third from the assumption that $\tau < c$.

and only if the resulting profit is non-negative: $p^* \frac{c-p^*+\tau}{2\tau} - c \frac{\beta-p^*}{\tau} \geq 0$. This inequality is strictly satisfied precisely for market clearing prices above \bar{p} (defined in the proposition), and in particular for our p^* . We have reached the desired state $\omega^* = (p^*, q^*, p^*, q^*)$. Note that all the transitions from one absorbing state to the other take place with the help of a single suitable experimentation.

Secondly, let us suppose that the initial state is of the form $\omega^0 = (p, q, p, q)$, where $p \in (\beta - \frac{\tau}{2}, \bar{p})$. In this state, the firms are not competing and they have their independent markets. If both firms are receiving negative profit, they shut down, giving rise to the absorbing state $(p, 0, p, 0)$. Irrespective of which of the two states is absorbing, let one firm experiment with the market clearing quantity $\frac{\beta-p}{\tau}$. The experimenting firm will receive the larger profit and gets immediately imitated and the process transits to the absorbing state $\omega^1 = (p, \frac{\beta-p}{\tau}, p, \frac{\beta-p}{\tau})$. Next, let one of the firms experiment with marginal cost pricing. Then, the experimenting firm receives zero profit, while its rival receives a negative profit (this is, as shown in *Part 1*, precisely what the condition $p < \bar{p}$ implies). As a result, marginal cost pricing gets imitated and we reach the state $\omega^2 = (c, \frac{\beta-p}{\tau}, c, \frac{\beta-p}{\tau})$. The sequence of single experimentations from ω^2 to a state in Ω^* has been illustrated in the previous paragraph.

Thirdly, suppose that $\omega^0 = (p, q, p, q)$ with $p \in (\bar{p}, \beta)$ but $q \neq \frac{\beta-p}{\tau}$. Obviously an experimentation to the market clearing quantity $\frac{\beta-p}{\tau}$ is imitated and we get to a state Ω^* with a single experimentation.

Finally, from state $\omega^0 = (p, 0, p, 0)$ with $p \geq \beta$, an experimentation to marginal cost is imitated with positive probability. The sequence of single experimentations from the resulting state to a state in Ω^* has been illustrated earlier.

We conclude that it is possible to reach a state in Ω^* from any other state not belonging to Ω^* with a sequence of single experimentations. Hence, the modified co-radius of Ω^* is 1.

Part 3. We now demonstrate that all states in Ω^* are stochastically stable. Take any two states ω' and ω'' in Ω^* . From ω'' we can go to the monopoly state $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau}, \frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ with 2 experimentations (one firm experiments to monopoly price and quantity; the other firm imitates as it now has a lower profit). The monopoly state does not belong to Ω^* and it follows from Part 2 that it is possible to transit further to ω' using a series of single experimentations. Hence, the “effective” resistance of transition from ω'' to ω' is equal to the radius of ω' , $R(\omega')$ (see Part 1), both of which are equal to 2. From Theorem 3 in Ellison (2000), it follows that if ω'' is stochastically stable, then so is ω' . Since, this is true for any two states ω' and ω'' in Ω^* , all states in Ω^* are stochastically stable. ■

Notice that all stochastically stable states imply market clearing and spatially segregated

markets (as $\bar{p} \geq \frac{\beta+c}{2} > \beta - \frac{\tau}{2}$). The price \bar{p} in combination with quantity $\bar{q} = \frac{\beta-\bar{p}}{\tau}$ results in zero profit for a firm if the rival prices according to marginal cost. The stochastically stable states are thus precisely those states that are immune against experimentations to marginal cost pricing. The value of \bar{p} , and hence the set of stochastically stable states, depends on the consumer valuation β . As \bar{p} equals the monopoly price at the lowest permissible value of β (while respecting the condition stated in the proposition) and is increasing in β , the stochastically stable states are comprised of prices higher than the monopoly price. So, where most other papers based on imitation learning lead to prices that are more competitive than the Nash prediction, we see that prices can also be higher than the Nash benchmark. However, in spite of the higher price, profits are lower than in the Nash equilibrium.

The condition $\beta < c + \tau$ implies that one firm is not able to attract the full mass of consumers with marginal cost pricing. This provides both firms with total control over part of the consumers: the consumers that never consider buying the rival's product. For this parameter range, in the monopoly state, it is possible for the rival firm to experiment to a price such that consumers are poached away to such an extent that the rival ends up with a higher profit. In a stochastically stable state, firms resort to catering to the set of consumers for which it is immune against any price experimentation by the rival firm. This results in the firms producing below monopoly quantity and charging a price above monopoly level. It is precisely this set of "high prices" (higher than even the monopoly price) which not only gives each firm independent markets, but under the restriction that pricing below marginal cost is not permissible, makes it impossible (evaluated in terms of relative profit) for the other firm to encroach on the other's turf. This gives rise to the emergence of niche markets that are stochastically stable.

Proposition 3.4. *Suppose $\beta \geq c + \tau$. Then, all absorbing states are stochastically stable.*

Proof. For the current parameter conditions, every absorbing state can be disrupted by an experimentation with marginal cost. Part 2 of the proof of Proposition 3.3 shows this for the states where the price is less than $\beta - \frac{\tau}{2}$. For the states where the price is larger than or equal to $\beta - \frac{\tau}{2}$, an experimentation to marginal cost yields a profit of 0 for the experimenting firm and leaves the other firm with a profit of $p(\frac{c-p+\tau}{2\tau}) - c\frac{\beta-p}{\tau}$. The latter profit is decreasing in β and is therefore maximized at $\beta = c + \tau$. At $\beta = c + \tau$ this profit is negative if and only if $p < c + \tau$ or $p > 2c$. So, for prices below $c + \tau$ this profit is negative. For prices above $c + \tau$, experimentation to marginal cost works trivially as it draws away all consumers.

Next, it is possible to reach the state $(c, \frac{1}{2}, c, \frac{1}{2})$ by a quantity experimentation. So, a $(c, \frac{1}{2}, c, \frac{1}{2})$ -tree with tree-resistance $|\Omega| - 1$ is easily constructed, where $|\Omega|$ is the cardinality

of the set of absorbing states. Since it requires at least one experimentation to leave any absorbing state, this is the least resistance any tree can have. In what follows, we show that for any absorbing state, a tree with resistance $|\Omega| - 1$ can be constructed and so, all absorbing states are stochastically stable.

First, consider the states (c, q, c, q) , where $q < \frac{1}{2}$. All that needs to be done in this case is to reverse the (c, q, c, q) to $(c, \frac{1}{2}, c, \frac{1}{2})$ edge in the just constructed $(c, \frac{1}{2}, c, \frac{1}{2})$ -tree and deleting all edges starting from (c, q, c, q) . The modified tree has the same resistance. Thus we have a (c, q, c, q) -tree, where $q < \frac{1}{2}$, with resistance $|\Omega| - 1$. Note that, in particular, the $(c, 0, c, 0)$ -tree has resistance $|\Omega| - 1$ and is stochastically stable.

Second, consider states of the form $(p, 0, p, 0)$. We can modify the just constructed $(c, 0, c, 0)$ -tree by reversing the $(p, 0, p, 0)$ to $(c, 0, c, 0)$ edge and deleting all edges starting from $(p, 0, p, 0)$. This leaves the resistance of the tree unchanged. Hence, states of that form are stochastically stable.

Third, consider states of the type (p, q, p, q) , where $p \neq c$ and q is such that both firms receive at least zero profit. To construct the (p, q, p, q) -tree we modify the corresponding $(p, 0, p, 0)$ -tree by deleting the edges starting from (p, q, p, q) and building an edge from $(p, 0, p, 0)$ to (p, q, p, q) . The constructed tree has a resistance of $|\Omega| - 1$. ■

In Figure 3.1 the set of stochastically stable states mentioned in the proposition is captured by the thick solid lines. In the states with prices equal to marginal cost or quantity equal to zero, the firms receive zero profit. Also the states at the curved part of the boundary imply zero profit, as this curve is composed of price–quantity pairs for which $p \cdot \frac{\beta-p}{\tau} - c \cdot q$ equals zero. In all remaining absorbing states, the firms experience a positive profit. Unlike

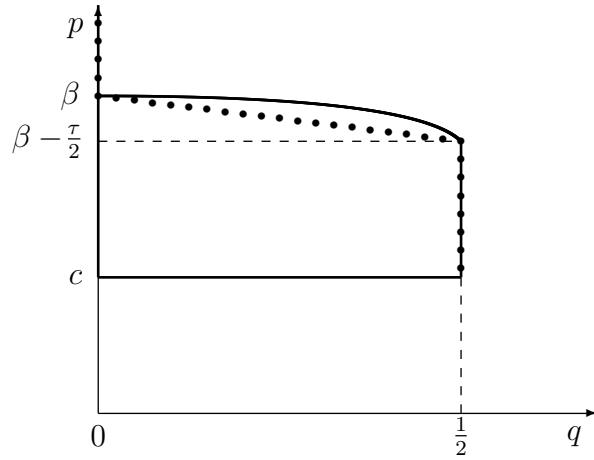


Figure 3.1: Illustration of the set of stochastically stable states.

in the previous two propositions, now there exist stochastically stable states where firms actively compete for the marginal agent. This is the case in all states with prices below

$\beta - \frac{\tau}{2}$. Also, unlike in the previous propositions, markets may not clear. Only in the states along the dotted path, markets clear. In all states above this path the market operates with excess supply, whereas below this path there is excess demand.

The huge multiplicity in stochastic stable states is mainly driven by our imitation dynamic that allows an active firm (positive quantity) with zero profit to imitate an inactive firm (zero quantity). Another plausible qualification that might be reasonable to consider is the imitation of active firms only, in line with the Activity Principle of Alós-Ferrer et al. (2000). Then starting from a zero profit absorbing state with positive production, such as $(c, \frac{1}{2}, c, \frac{1}{2})$, experimentation by one firm to zero production will not induce a transition (with positive probability) to the state $(c, 0, c, 0)$ in the unperturbed (imitation) dynamic. With this restriction in place, only the absorbing states with an arbitrary price and no production – i.e. $(p, 0, p, 0)$ – are now excluded from being stochastically stable. The refinement that is obtained is, therefore, not qualitatively substantive.

An alternative assumption on the imitation dynamic is that while a firm imitates the other firm when the latter has strictly higher profit, in case of equal profit, only the firm with larger sales is imitated. This is equivalent to firms having lexicographic preferences where they primarily care about relative profit and secondarily about sales. Then, $(c, \frac{1}{2}, c, \frac{1}{2})$ will be the unique stochastically stable state. In particular, all states with excess demand or supply will no longer be stochastically stable. Remarkably, the only stochastically stable state implies prices to be equal to marginal costs, in spite of the differentiated market.

3.4 Discussion

Proposition 3.4 reveals that excess supply or excess demand can be a persistent feature of a horizontally differentiated market where firm behavior is guided primarily by imitation of the better off firm. Two natural questions arise: why do firms not cut down on production when there are unsold goods? and why do firms not produce more when there is unserved demand? If we were to say that a firm adjusts accordingly on the observance of non-market clearing, then we add something apart from imitation to our framework. By consideration of such adjustments, the non-market clearing states would no longer be stochastically stable simply because they cease to be absorbing states. Even then, the result that all absorbing states are stochastically stable would hold – the difference of course being that all such states are market clearing.

Next, let us comment on the robustness of the findings with respect to alternative assumptions and consider possible extensions to the model. Firstly, we could assume that firms have a status-quo bias and they do not imitate the rival when they receive equal profit with different strategies; in this case, all states where firms receive equal profit (possibly

with disparate strategies) would become absorbing under the unperturbed imitation dynamic. The set of stochastically stable states in Proposition 3.4 would expand accordingly to include these absorbing states, but this modification would leave Propositions 3.2 and 3.3 unaltered.

Secondly, let us consider the implications of our assumption of differentiated products. A perfectly homogenous market is a limiting case of our model (where zero transportation cost equals zero). It is easily seen that stochastic stability loses all its refinement power and all absorbing states will be stochastically stable. After all, from any state where the price is different from the marginal cost, a firm experimenting to marginal cost pricing will attract all consumers who are willing to buy at marginal cost price. This yields zero profit for the experimenting firm, but this is at least as much the profit of the rival firm, which faces zero demand and so may make negative profits if it produces a positive quantity. With positive probability the rival will imitate and, hence, it takes a single experimentation to reach a state where the price is equal to marginal cost. Next, by an argument similar to the one made in the proof of Proposition 3.4, it is possible to construct a path from any absorbing state to another in which the edges have a resistance of one and so, stochastic stability cannot refine the set of absorbing states.

Finally, we take up the issue of observability of decisions. In our model, we assumed prices, quantities and profits to be perfectly observable. In some situations, it may appear more realistic to assume that, rather than the quantity produced, the quantity sold is observable. There is a disparity between observation of quantity produced and observation of sales only when the quantity produced is in excess of the sales. This in turn can have an effect on the process only if the more successful firm produces in excess of its demand. Hence, observability of sales can only have a differential effect on Proposition 3.4, as it is the only situation where states with excess demand are stochastically stable. In Proposition 3.4, at a certain price p , the maximum quantity that can be produced in a stochastically stable state is $q = p\frac{\beta-p}{c\tau}$ – the quantity that yields zero profit, according to the equation: $p(\frac{\beta-p}{\tau}) - cq = 0$. But at a price p , the sales of a firm will be at most $\frac{\beta-p}{\tau}$; so, with observability of sales, any quantity greater than $\frac{\beta-p}{\tau}$ will not be imitated. Hence, absorbing states with price $p > \beta - \frac{\tau}{2}$ and quantity in the interval $(\frac{\beta-p}{\tau}, p\frac{\beta-p}{c\tau}]$ that were stochastically stable earlier, will cease to be so. On the other hand, in case the price $p \leq \beta - \frac{\tau}{2}$ and production $q > \frac{1}{2}$, both firms observe the rival receiving equal profit by selling $\frac{1}{2}$. Then, with positive probability, they imitate the rival's sales. As a result the states where $p \leq \beta - \frac{\tau}{2}$ and production $q > \frac{1}{2}$, which represented excess supply, cease to be absorbing states. Thus, observability of sales rules out excess supply as possible property of a stochastically stable state. The main thrust of Proposition 3.4 that there exist stochastically stable states that are not market clearing remains unaltered, though now only with excess demand.

3.5 Conclusion

In this chapter, we examined the stochastically stable outcome when two firms in a differentiated market choose price and quantity, and market clearing is not assumed a priori. Firms adapt their choice by imitating the rival firm if the rival has a higher (relative) profit. For low values of consumer valuation of the product, in the stochastically stable state, the firms choose as if they were monopolists, i.e. the price and quantity that maximizes profit in absence of a rival. For intermediate values of consumer valuation, the stochastically stable state consists of a set of prices higher than the monopoly price and the corresponding market-clearing quantities. Finally, for large values of consumer valuation, all absorbing states are stochastically stable; so, non-market clearing states may also be stochastically stable. In terms of the profit that firms receive in the stochastically stable state, when the consumer valuation is low, they extract maximal surplus from the market (this also coincides with the profit level of the corresponding Nash equilibrium); for intermediate values of consumer valuation, because prices are higher, the firms receive lower profit compared to the profit at corresponding Nash equilibrium. However, it is not possible to make a similar clear-cut comparison when the consumer valuation is large.

An interesting follow-up question would be to study this imitation dynamic while allowing for a larger number of firms, which might then permit interesting comparative static insights. To elaborate, we have argued that the stochastically stable states under the specified imitation dynamic are the ones that are most resistant to any form of spiteful behaviour; if we allow for more firms, we speculate that successful spiteful behaviour would become more difficult as it has to be strong enough to affect a larger number of firms. As a result, on the one hand, it may be possible that the stochastically stable states would yield larger profits when there are more firms. However, at the same time, as the number of number of firms increase, the market power of each individual firm decreases and this might have a negative effect on the profit levels in the stochastically stable states. There is thus, the possibility of a conflict between these two forces and it might be instructive to examine the interplay between these two opposing tendencies. This subject is examined in the next chapter.

3.6 Appendix

We present the stochastically stable states for the case when transportation costs are larger than the marginal cost of production (i.e. $\tau \geq c$). The results are segregated by the consumer valuation, β .

Proposition 3.5. *Suppose $\beta < (\tau - 2c + \sqrt{8c^2})$. The state where both firms supply a quantity of $q^* = \frac{\beta-c}{2\tau}$ and set a price equal to $p^* = \frac{\beta+c}{2}$ is the unique stochastically stable state.*

Proof. Let us denote the state mentioned in the proposition by ω^* . The supposition in the proposition implies $\beta < c + \tau$, and so, the quantity produced by each firm according to ω^* is less than half the market size. Our proof is constructed as follows. First, we show that no experimentation in price *and* quantity can push the system outside the basin of attraction of state ω^* . It follows that it is then also not possible with *either* a price *or* a quantity experimentation. From this we conclude that the radius of ω^* is strictly larger than 2. Second, we show that the co-radius of ω^* is at most equal to 2. From part (i) of Theorem 2 in Ellison (2000), it follows that all stochastically stable states are included in the set $\{\omega^*\}$. Next, existence leads us to conclude that ω^* is the unique stochastically stable state.

Part 1 (radius). In state ω^* the firms set price and quantity at the monopoly level and do not compete for the marginal agent. Experimentation in price by firm 0 leads to active competition with firm 1 if and only if the experimented price is less than $\frac{3}{2}\beta - \frac{1}{2}c - \tau$. Notice that any larger price leads to a decrease in firm 0's profit while it does not affect firm 1's profit and will result in a return to ω^* (firm 0 will imitate firm 1). Suppose, therefore that firm 0 experiments with the price $p' < \frac{3}{2}\beta - \frac{1}{2}c - \tau$. Moreover, suppose firm 0 experiments simultaneously with the optimal quantity given that price: $q' = \frac{\beta+c+2\tau-2p'}{4\tau}$. This experimentation leads to profits of

$$\pi'_0 = (p' - c)q' \quad \text{and} \quad \pi'_1 = \frac{\beta+c}{2}(1 - q') - cq^*$$

for firm 0 and firm 1 respectively. The difference between these two profits, $\pi'_0 - \pi'_1$, is maximized at $p'_{max} = \frac{c+\tau}{2}$. If firm 0 experiments to $p' = \frac{c+\tau}{2}$ and $q' = \frac{\beta+\tau}{4\tau}$. Then, the difference in profit between the two firms, $\pi'_0 - \pi'_1$, simplifies to

$$\frac{1}{8\tau}(\beta^2 - 2\tau\beta + 4c\beta + \tau^2 - 4c^2 - 4c\tau).$$

The basin of attraction of state ω^* cannot be left if and only if this experimentation by firm 0 leads to a situation where firm 1's profit is not less than that of firm 0. This is precisely the

case when the expression above is non-positive i.e. when $\beta \in [\tau - 2c - \sqrt{8c^2}, \tau - 2c + \sqrt{8c^2}]$. Now, notice that $\beta < \tau - 2c - \sqrt{8c^2}$ implies $p'_{max} = \frac{c+\tau}{2} > \frac{\beta+c}{2}$, which is inconsistent with the experimented price being below $\frac{3}{2}\beta - \frac{1}{2}c - \tau$. (This conclusion follows from the fact that $\beta < c + \tau$ implies $\frac{3}{2}\beta - \frac{1}{2}c - \tau < \frac{\beta+c}{2}$.) Therefore, in this case, we have to look at the experimentation with $p' = \frac{3}{2}\beta - \frac{1}{2}c - \tau$, and it can easily be verified that with this experimentation, the relative profit expression is negative; so, there cannot be a transition from the basin of attraction of state ω^* .

We have shown that even in the best case experimentation with prices and quantities, firm 0 cannot obtain a situation where its profit strictly exceeds that of firm 1 when the conditions in the proposition are met. Now, note that if it is not possible to leave ω^* with a price experimentation in combination with an optimal price adjustment, it is also not possible without such an adjustment. A sole quantity experimentation always decreases the profit (either because more units can be sold at a price above cost, or because more units are produced than are sold), while it doesn't affect the profit of the rival firm.

Part 2 (co-radius). From Part 1, we learn that starting from any situation, if firm 0 experiments simultaneously to price p^* and quantity q^* , then it results in a larger profit relative to the profit of firm 1. This means that, from any state, at most two experimentations are needed to get to (the basin of attraction of) ω^* . ■

Proposition 3.6. Suppose $\tau - 2c + \sqrt{8c^2} \leq \beta < \frac{3}{2}\tau$.⁹ The set of stochastically stable states consists of the monomorphic states with prices $p^* \in [\beta - \frac{\tau}{2}, \hat{p}]$ and quantity $q^* = \frac{\beta-p^*}{\tau}$, where $\hat{p} = \min\{c + \tau, \frac{3c+\tau+\sqrt{9c^2+\tau^2+6c\tau-8c\beta}}{2}\}$.

Proof. Let Ω denote the set of absorbing states and Ω^* denote the set of stochastically stable states as claimed in the proposition. By construction, $\hat{p} \leq c + \tau$. Notice that states contained in Ω^* imply market clearing, and in none of them firms actively compete for the marginal consumer. Our proof is constructed as follows. First, we show that the modified co-radius of Ω^* is 1. Second, we show that the radius of Ω^* is 2. It follows from Ellison (2000) that the set of stochastically stable states is contained in Ω^* . Third, we show that all states in Ω^* are stochastically stable. This concludes the proof.

Part 1 (modified co-radius) We demonstrate that the modified co-radius of the set Ω^* is 1, by showing that from all other states in Ω , it is possible to transit to a state in $\hat{\Omega}$ with a series of experimentations in either price or quantity only.¹⁰

⁹For some values of τ , with $\tau \geq c$, it is possible that the parameter restriction is not applicable, in which case one has to consider the Proposition 3.7.

¹⁰Note that in the proof we ignore the consideration that the firms have negative profit and shut down. In case of a shut down at a price greater than or equal to $\beta - \frac{\tau}{2}$ experimentation to a quantity of $\frac{\beta-p}{\tau}$ is imitated; for lower prices a quantity experimentation to $\frac{1}{2}$ is imitated.

First, assume we are in a state of $\Omega \setminus \Omega^*$ with $p < \beta - \frac{\tau}{2}$. In case $q \neq \frac{1}{2}$, an experimentation with $q = \frac{1}{2}$ leads to at least as much profit for the experimenting firm while not changing the profit of the competing firm. Due to the positive probability of imitation, a single experimentation is needed to get to an absorbing state with $p < \beta - \frac{\tau}{2}$ and $q = \frac{1}{2}$. From here, a price experimentation with $p' = \beta - \frac{\tau}{2}$ yields the experimenting firm a profit of $\pi' = (\beta - \frac{\tau}{2})\frac{1}{2\tau}(p - \beta + \frac{3}{2}\tau) - \frac{1}{2}c$, while its rival stays at a profit of $\pi = \frac{1}{2}p - \frac{1}{2}c$. The experimenting firm has a larger profit, as by $\beta < \frac{3\tau}{2}$,

$$\pi' > \pi \iff (\beta - \frac{3}{2}\tau)p > (\beta - \frac{3}{2}\tau)(\beta - \frac{\tau}{2}) \iff p < \beta - \frac{\tau}{2}.$$

Hence, a single experimentation is needed to transit to the state where both firms set the price at $p = \beta - \frac{\tau}{2}$ and the quantity at $q = \frac{1}{2}$, which is a state in Ω^* .

Second, assume we are in a state of $\Omega \setminus \Omega^*$ with $p \geq \hat{p} > \beta - \frac{\tau}{2}$. If $q \neq \frac{\beta-p}{\tau}$, then an experimentation by one firm to this quantity will be immediately imitated. So, we focus on initial absorbing states in $\Omega \setminus \Omega^*$, where $p \geq \hat{p} > \beta - \frac{\tau}{2}$ and $q = \frac{\beta-p}{\tau}$. Now, consider an experimentation by one of the firms to marginal costs. Clearly, the experimenting firm ends up with zero profit. By construction, \hat{p} can take a value of at most $c + \tau$. If $\hat{p} = c + \tau$, then $p > c + \tau$, and the non-experimenting firm will see all its consumers disappear and will make a zero profit at best (and in fact, negative in case of production of a positive quantity). On the other hand, if $p < c + \tau$ (which is only possible if $\hat{p} < c + \tau$), the non-experimenting firm will attract $\frac{c-p+\tau}{2\tau}$ consumers and receive a profit of $p\frac{c-p+\tau}{2\tau} - c\frac{\beta-p}{\tau}$. This profit is negative if and only if $p^2 - (3c + \tau)p + 2c\beta > 0$, which is satisfied if $p > \frac{3c+\tau+\sqrt{9c^2+\tau^2+6c\tau-8c\beta}}{2} = \hat{p}$ (the equality follows from the construction of \hat{p} , and the assumption that $\hat{p} < c + \tau$). In any of the two cases, with positive probability the experimenting firm is imitated and a state is reached where both firms price at marginal cost. From this absorbing state where the price equals c (and since $c < \beta - \frac{\tau}{2}$), we can invoke the arguments in the previous paragraphs to realise that it takes a series of single price or quantity experimentations to move into a state in Ω^* .

Finally, from the absorbing states with price $p \in [\beta - \frac{\tau}{2}, \hat{p}]$ and $q \neq \frac{\beta-p}{\tau}$ only a single quantity mutation to $q = \frac{\beta-p}{\tau}$ is needed to get to a state in Ω^* . Thus, we have shown that it takes a series of single price or quantity experimentations to move into a state in Ω^* , from any other absorbing state.

Part 2 (radius). It is not possible to move out of any state in Ω^* with a single experimentation in price or quantity. Since in all these states markets clear while the markets are segregated and prices are above marginal cost, only quantity experimentations would leave the experimenting firm strictly worse off. For this reason, only price experimentations have to be considered. Further, it can be shown (and is in fact shown in the proof of Proposition 3.7) that it suffices to examine the experimentation by a firm to marginal cost,

the reason for this being that if an absorbing state is immune to a marginal cost pricing experimentation, then it is immune from all price experimentations. It can be verified that the absorbing state in Ω^* where both firms price at $\beta - \frac{\tau}{2}$ is immune against marginal cost experimentation, as is the state where both firms price below \hat{p} . By concavity of the (relative) profit function it follows that all states with prices larger than and equal to $\beta - \frac{\tau}{2}$ and less than \hat{p} are immune to marginal cost experimentations. We can conclude that the radius is at least 2.

Next, we show that the radius is precisely 2 by showing that states in Ω^* can be disrupted by a firm making a price and quantity experimentation. If $p \notin \{\beta - \frac{\tau}{2}, \frac{\beta+c}{2}\}$, experimentation to the price-quantity pair $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ in case $\beta \leq c + \tau$ and to $(\beta - \frac{\tau}{2}, \frac{1}{2})$ in case $\beta > c + \tau$ yields a higher profit to the experimenting firm relative to its rival. The experimenting firm will be imitated resulting in a transition to a market-clearing absorbing state where, depending on the relative value of β and $c + \tau$, the price is $\beta - \frac{\tau}{2}$ or $\frac{\beta+c}{2}$. Now from this absorbing state, a price and quantity experimentation to $(\frac{c+\tau}{2}, \frac{1}{2\tau}(p - \frac{c+\tau}{2} + \tau))$ results in a transition from this state, and from Ω^* .¹¹

Part 3. We now show that all states in Ω^* are in the set of stochastically stable states by showing that $R(\omega') = c^*(\omega'', \omega') = 2$, for any two states ω'' and ω' in Ω^* .¹²

First suppose that $\beta < c + \tau$. If ω'' is the absorbing state where firms choose price $\frac{\beta+c}{2}$ (with quantity $\frac{\beta-c}{2\tau}$), then as mentioned in Part 2 above, a price and quantity experimentation to $(\frac{c+\tau}{2}, \frac{1}{2\tau}(p - \frac{c+\tau}{2} + \tau))$ results in a transition from this state, and the new absorbing state is outside Ω^* .¹³ From here, using the arguments of Part 1 above, it is possible to transit to any other state ω' in Ω^* with either price or quantity experimentations. So, in this case, $c^*(\omega'', \omega') = 2$. On the other hand, if in the absorbing state ω'' , firms choose price other than $\frac{\beta+c}{2}$, then it takes a price and quantity experimentation with $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$, to reach the state where both firms choose $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$. We have just outlined the process thereafter, and so, even in this case, $c^*(\omega'', \omega') = 2$. We already know from Part 2 that $R(\omega') = 2$, for any state $\omega \in \Omega^*$. Thus, when $\beta < c + \tau$, for any two states ω'' and ω' in Ω^* , $R(\omega') = c^*(\omega'', \omega') = 2$.

Finally, if $\beta \geq c + \tau$, then we repeat the arguments made in the preceding paragraph with the only change being that the state where firms choose price $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ is replaced by the state where firms choose price and quantity $(\beta - \frac{\tau}{2}, \frac{1}{2})$. Amongst all absorbing states

¹¹It is easily verified that $\frac{c+\tau}{2} < \beta - \frac{\tau}{2}$, and so the state $(\frac{c+\tau}{2}, \frac{1}{2\tau}(p - \frac{c+\tau}{2} + \tau))$ is not in the set Ω^* .

¹² $c^*(\omega'', \omega')$ denotes the effective resistance between the absorbing states ω'' and ω' . For the effective resistance, the radius of all intermediary absorbing states is subtracted from the total resistance of the transition from ω'' to ω' .

¹³When $\beta < c + \tau$, then $\frac{\beta+c}{2} > \beta - \frac{\tau}{2}$, and so the state where firms choose price $\frac{\beta+c}{2}$ and quantity $\frac{\beta-c}{2\tau}$ belongs to Ω^* . Further, as this state represents the monopoly absorbing state, firms get the highest profit in this state.

in Ω^* , this state now represents the absorbing state that yields the highest profit for each firm. Thus, when $\beta \geq c + \tau$, for any two states ω'' and ω' in Ω^* , $R(\omega') = c^*(\omega'', \omega') = 2$.

It follows from Ellison (2000), that all absorbing states in Ω^* belong to the set of stochastically stable states. ■

Proposition 3.7. Suppose $\beta \geq \max\{\tau - 2c + \sqrt{8c^2}, \frac{3}{2}\tau\}$. Let $\hat{\beta} = \min\{\tau + \frac{3c}{2}, \frac{9c^2 + \tau^2 + 6c\tau}{8c}\}$. If $\beta < \hat{\beta}$, then the set of stochastically stable states consists of the states with prices $p^* \in (\underline{p}, \bar{p})$ and quantity $q^* = \frac{\beta - p^*}{\tau}$, where $\underline{p} = \max\{\frac{3c + \tau - \sqrt{9c^2 + \tau^2 + 6c\tau - 8c\beta}}{2}, \beta - \frac{\tau}{2}\}$ and $\bar{p} = \min\{\frac{3c + \tau + \sqrt{9c^2 + \tau^2 + 6c\tau - 8c\beta}}{2}, c + \tau\}$. If $\beta \geq \hat{\beta}$, then all absorbing states are stochastically stable.

Proof. Let Ω^* denote the set of stochastically stable states as claimed in the first part of the proposition. Notice that all states in Ω^* imply market clearing and in none of them firms actively compete for the marginal consumer ($p > \underline{p} \geq \beta - \frac{\tau}{2}$). Our proof for the first part is constructed as follows. First, we show that the radius of Ω^* is 2. Second, we show that the modified co-radius of Ω^* is 1. Next, it follows from Ellison (2000) that the stochastically stable states are contained in Ω^* . Third, we show that all states in Ω^* are stochastically stable. That concludes the proof.

Part 1 (radius). We demonstrate that the radius of the set Ω^* is 2. We first show that the radius is at least 2 by showing that there exists no price experimentation or quantity experimentation that makes the system leave the basin of attraction of Ω^* .

Since supply equals demand in all states in Ω^* and the consumers' acquisition decision is independent of the quantity produced, a quantity experimentation reduces the own profit (given that the price is greater than the marginal cost) while it leaves the rival's profit unaffected. Hence, a quantity experimentation leaves the experimenting firm worse off as compared to its rival and imitation will drive the process back to the original state.

Since firms do not compete for the marginal agent, experimentation to a larger price ($p' > p$) will not lead to a change in the rival's profit. Moreover, since the quantity is unadjusted, the cost component of the own profit does not change. The revenue component increases if $p' \frac{\beta - p'}{\tau} > p \frac{\beta - p}{\tau}$, which is satisfied if and only if $p' \in (\beta - p, p)$; a contradiction.¹⁴ Experimentation to a lower price ($p' < p$) decreases the own profit, and can therefore only lead to a transit to another absorbing state if the rival's firm is decreased by even more. For this to be the case, the price p' must be sufficiently low to compete for the rival's consumers: $p' < 2\beta - p - \tau$. Such an experimentation results in a situation where the experimenting firm obtains a larger profit than its rival if and only if it leads to a larger revenue (since costs are equal): $p' \frac{\beta - p}{\tau} - p' \frac{-p + \tau}{2\tau} \geq 0$. Since the derivative of the left-hand

¹⁴Note that $\beta - p < p$, since $p > \underline{p} \geq \beta - \frac{\tau}{2} \geq \frac{\beta}{2} + \frac{3}{4}\tau - \frac{1}{2}\tau > \frac{\beta}{2}$.

side to p' is constant ($2\beta - 3p$), it suffices to check the validity of the inequality at the two extreme values: c and $2\beta - p - \tau$. Experimentation to $2\beta - p - \tau$ can not be disruptive for the obvious reason that no consumers are stolen away and the own profit decreases. For $p' = c$ the inequality simplifies to $p^2 - (3c + \tau)p + 2\beta c \geq 0$. Contrariwise, a necessary condition for the price p being immune against marginal cost experimentation is $p^2 - (3c + \tau)p + 2\beta c < 0$; that is, $p > \frac{3c + \tau - \sqrt{9c^2 + \tau^2 + 6c\tau - 8c\beta}}{2}$ and $p < \frac{3c + \tau + \sqrt{9c^2 + \tau^2 + 6c\tau - 8c\beta}}{2}$. The remaining necessary conditions are that $p > \beta - \frac{\tau}{2}$ (prices below this cannot withstand a price deviation; see Part 2 for more explicit reasoning on this) and $p < c + \tau$ (otherwise, marginal cost pricing takes away all clientele). These conditions give rise to the interval (\underline{p}, \bar{p}) .

The lower bound on the values for β , come from comes from Proposition 3.5 ($\beta \geq \tau - 2c + \sqrt{8c^2}$) and Proposition 3.6(i) ($\beta \geq \frac{3}{2}\tau$). The requirement $\beta < \frac{9c^2 + \tau^2 + 6c\tau}{8c}$ makes sure that the roots in the previous paragraph are well-defined (if not, then the statement is void). Finally, the price of the rival for which experimentation to marginal cost would lead to the largest profit difference in favor of the experimenting firm is $p = \frac{3c + \tau}{2}$. This price must exceed $\beta - \frac{\tau}{2}$ to ensure that not all absorbing states are defeated by marginal cost experimentation (see Part 2 below). Therefore, $\beta < \tau + \frac{3c}{2}$.

We can conclude that for the states in Ω^* , the radius is at least 2. Since all states in Ω^* imply segregated markets, an experimentation by one firm to monopoly price *and* quantity leads to an immediate drift towards the monopoly state (if not there yet). In case the monopoly state is not in Ω^* we can conclude already that the radius of Ω^* is precisely 2. Otherwise, from the monopoly state another experimentation in price *and* quantity to $p' = \frac{c+\tau}{2}$ and $q' = \frac{\beta-\tau}{4}$) will induce a transition outside Ω^* as $p' < \beta - \frac{\tau}{2}$. Again, we can conclude that the radius of Ω^* is precisely 2.

Part 2 (modified co-radius). We now show that the modified co-radius of Ω^* is 1. We do so by demonstrating that from any other state outside Ω^* , it is possible to reach any state in Ω^* with a sequence of single suitable experimentations. This implies that the modified co-radius of Ω^* is 1.

Firstly, let us suppose that the initial state is of the form $\omega^0 = (p, q, p, q)$ with $p > \beta - \frac{\tau}{2}$ but $p \notin (\underline{p}, \bar{p})$. In this state, firms do not compete for the marginal consumer. And, let $\omega^* = (p^*, q^*, p^*, q^*)$ be the state in Ω^* that we try to reach. If $q \neq \frac{\beta-p}{\tau}$, then it takes one quantity experimentation to move to the state $\omega^1 = (p, \frac{\beta-p}{\tau}, p, \frac{\beta-p}{\tau})$. From this state only one price experimentation to marginal cost level is needed to transit to $\omega^2 = (c, \frac{\beta-p}{\tau}, c, \frac{\beta-p}{\tau})$.¹⁵ In this state both firms receive zero profit and experimentation with quantity q^* allows a transit to $\omega^3 = (c, q^*, c, q^*)$. Now a mutation to the corresponding price p^* makes the system move to $\omega^* \in \Omega^*$.

¹⁵Note that Ω^* is defined to consist of states that are immune to marginal cost experimentation.

Secondly, consider the state $\omega^0 = (p, q, p, q)$ with $p \leq \beta - \frac{\tau}{2}$. If $q \neq \frac{1}{2}$, one quantity experimentation is needed to get to $\omega^1 = (p, \frac{1}{2}, p, \frac{1}{2})$. Next, consider a price experimentation to τ . In case $\tau > p$, the experimenting firm receives a revenue of $\tau \frac{p}{2\tau}$ which is equal to that of and its rival $p \frac{1}{2}$. In case $\tau < p$, the experimenting firm receives a revenue of $\tau \frac{1}{2}$ which is larger than the revenue of its rival $p \frac{2\tau-p}{2\tau}$. As in both cases the costs are equal ($\frac{1}{2}c$), the system can transit to $\omega^2 = (\tau, \frac{1}{2}, \tau, \frac{1}{2})$. Now, it takes just a price experimentation to marginal cost to get to $\omega^3 = (c, \frac{1}{2}, c, \frac{1}{2})$. We have earlier illustrated the sequence of single mutations that takes the system to a state in Ω^* .

We have demonstrated that it is possible to transit to a state in Ω^* from all states outside Ω^* with a series of single experimentations. The modified co-radius of Ω^* is equal to 1. We then conclude that all the stochastically stable states are contained in Ω^* .

Part 3.

We now demonstrate that all states in Ω^* are stochastically stable. Take any two states ω' and ω'' in Ω^* .

First, assume that the monopoly state $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau}, \frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ does not belong to Ω^* . From ω'' we can go to the monopoly state $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau}, \frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ with 2 mutations (one firm experiments to monopoly price and quantity; the other follows). Next, only a price experimentation with marginal cost is needed to transits to a state where both firm price according to marginal cost. Part 1 illustrated how to go from here to the state $\omega' \in \Omega^*$ with sequence of single experimentations in price or quantity. It follows that $R(\omega') = c^*(\omega'', \omega') = 2$.

Second, assume that the monopoly state belongs to Ω^* . Now, it is not possible to leave the basin of the monopoly state with one single price experimentation. We modify the previous argument as follows: With 2 experimentations to move from ω'' to the monopoly state. (If ω'' is the monopoly state, then we start the process from hereon. If ω' is the monopoly state, then we stop the process at this point.) From the monopoly state, let one firm experiment in both price and quantity: $p = \frac{c+\tau}{2}$ and $q = \frac{\beta-\tau}{4}$. The experimenting firm will receive at least as much profit as its rival and, as a consequence, we transit to the state (p, q, p, q) . Now, it is possible to transit to ω' via a series of experimentations in one dimension (see Part 1). Hence, $R(\omega') = c^*(\omega'', \omega') = 2$.

From Ellison (2000) it follows that if ω'' is stochastically stable, then so is ω' . Since, $R(\omega') = c^*(\omega'', \omega')$ for any two states ω' and ω'' in Ω^* , it follows that all states in Ω^* are stochastically stable. This proves the first part of the Proposition.

Now, we prove the second part of the proposition. Here the statement applies to the situation when $\beta \geq \max\{\hat{\beta}, \frac{3\tau}{2}\}$ with $\hat{\beta} = \min\{2\tau + \frac{c}{2}, \frac{9c^2 + \tau^2 + 6c\tau}{8c}\}$. In this situation there does not exist any state that is immune against price experimentations and from any state it is possible to reach the state $(c, \frac{1}{2}, c, \frac{1}{2})$ with a series of single experimentations. If the price is greater than $\beta - \frac{\tau}{2}$ this can be done in one step. Otherwise, two steps are needed: a

price experimentation to a price of τ (as in Part 2), followed by a price experimentation to marginal cost (also, as in Part 2). The stochastic potential of the least resistant $(c, \frac{1}{2}, c, \frac{1}{2})$ -tree is $|\Omega| - 1$, as all edges of this ω -tree have a weight of 1. This implies stochastic stability of this state. All absorbing states of the process are also stochastically stable, as it is possible to modify the above tree and construct a ω -tree of resistance $|\Omega| - 1$ for any absorbing state ω (described in more detail in Proposition 3.4). \blacksquare

Chapter 4

A Multi-firm Imitation model¹

4.1 Introduction

In the previous chapter, motivated by Alchian (1950), who proposes an alternate economic perspective that emphasises relative performance over absolute performance, we analysed the outcome of imitation driven strategy choice in a duopoly. In this chapter, we study the evolution of imitative behaviour in a (spatially) differentiated market. Firms are equidistantly located on a circle of unit circumference, while consumers are located uniformly along the circumference.² The extent of differentiation is reflected in the transportation cost that consumers incur when they travel to a firm to acquire the good. Firms choose both price and quantity, and market-clearing is not assumed a priori. This is descriptive of situations where production may be in-advance (as in a Cournot model and as opposed to a Bertrand model where production is on-demand) and yet firms have the independence to choose price (as in a Bertrand model but unlike a Cournot model). A part of our analysis concerns the possibility of “endogenous” market-clearing.³

The firms are assumed to operate in an environment where they may not be fully informed of demand related conditions. In particular, they need not know that they are in a differentiated market, the consumers’ valuation of the product or the consumers’ preferences between the firms. They are only aware of their own cost and the presence of other firms. The strategy of each firm and its profit is observable. The informational constraints may compel a firm to use a rule of thumb which, in this case, is specified by

¹This chapter is based on the paper “Imitation by Price and Quantity setting firms in a Differentiated Market”, Meteor Research Memorandum 13/022, and is co-authored with Ronald Peeters.

²While the assumption of equidistantly located firms is made merely to specify the demand structure, it is interesting to note that Hohenkamp and Wambach (2010) comment that this would also be the stochastically stable outcome of an evolutionary model where firms choose their location on the Salop circle.

³Herings (1997) and Alós-Ferrer and Kirchsteiger (2010) show that it is possible for non-clearing markets institutions to evolve.

imitation of the strategy of the firm with the highest profit. Occasionally a firm may experiment with other strategies, and our interest lies in the stochastically stable outcome, which is the outcome one would expect to emerge from the process of imitation as the tendency to experiment goes to zero.

The results indicate that the stochastically stable outcome of this process of imitation coincides with the (unique symmetric) Nash equilibrium of the underlying game. For high levels of differentiation, each firm charges a “monopoly” price and chooses to produce the demand that it faces at this price. Of course, the “monopoly” price that each firm chooses is independent of the number of firms and the level of differentiation. For intermediate levels of differentiation, each firm chooses a “mutually non-aggressive” price and produces the corresponding demand. Counter-intuitively, in this case, an increase in the number of firms increases the price that a firm charges (*ceteris paribus*). Additionally, an increase in the level of differentiation decreases the price charged by a firm (*ceteris paribus*). The implication of this is that more competition may actually increase prices while greater market differentiation may reduce them. Finally, for low levels of differentiation, each firm charges a mark-up over the marginal cost, where the mark-up varies positively with the level of product differentiation and negatively with the number of firms. The quantity chosen by each firm always equals the demand that it faces, so that the stochastically stable states are all endogenously market-clearing.

It is interesting to see the coincidence of the outcome of this “low-rationality” process with the Nash equilibrium, as it is often observed that an imitation protocol leads to spiteful behaviour that is detrimental for all involved. Due to the importance accorded to relative payoffs (or in this case, relative profit), a firm may be driven to adopt a strategy that while reducing its own payoff, reduces that of its rivals to a larger degree.⁴ This is illustrated in the evolutionary Cournot model in Vega-Redondo (1997), where firms choose quantities in a homogeneous market, and the total quantity produced determines the market price. In Vega-Redondo (1997), the quantity-choice of each firm and its profit is observable to all. A typical firm (almost always) imitates the firm that received the highest profit, while occasionally choosing a random quantity. The stochastically stable outcome of this process is the Walrasian equilibrium, which is not a Nash equilibrium outcome. Tanaka (1999) extends this to asymmetric oligopolies where two groups of firms – low cost and high cost – exist, and finds that the stochastically stable state under imitation corresponds to the

⁴Bergin and Bernhardt (2004) make the point in an imitation dynamic, as the comparative performance looms large, firms end up ignoring the externality effect of their own actions. This often leads to payoffs that are lower than Nash equilibrium payoffs, and such an occurrence was seen in the previous chapter. Here, in contrast, the stochastically stable state under imitation coincides with the Nash equilibrium and one of the driving factors is that competition amongst firms is primarily local (i.e. more intense with the neighbouring firm) while imitation is global.

Walrasian output in each group of firms.⁵ Even though imitation-driven spiteful incentives may play a role here as well, we find that the Nash outcomes are expected to emerge in the long-run, and we present the intuition for this at a later stage.

The framework and analysis in this chapter is complimentary to Tanaka (2000, 2001). Tanaka (2000) analyses a symmetric oligopoly in differentiated markets with a linear demand function, and finds that the stochastically stable behaviour under price or quantity imitation coincides with the unique finite-population evolutionarily stable strategy. Tanaka (2001) shows that this coincidence of the stochastically stable states under price imitation or quantity imitation carries over to differentiated markets with linear demand functions, even when firms are asymmetric with respect to the cost of production. In this chapter, we stress on a feature that we consider to be vital (in some situations) in differentiated markets: the feature of spatial differentiation. In differentiated markets, it is plausible that a particular firm is relatively more affected by the strategy of a subset of firms. This nuance is not captured in a differentiated market used in the earlier papers, but is conveniently encapsulated by the spatial differentiation structure of the Salop circle – a firm is more affected by, and in turn affects to a greater extent, the immediately neighbouring firms. We complement existing literature in evolution of firm behaviour by introducing this feature of asymmetry in differentiation, and by allowing firms to choose both price and quantity thereby making more explicit the process of price and quantity determination.

The chapter is structured as follows. In Section 4.2 we present the model. Section 4.3 contains the results. We begin by characterising the Nash equilibrium of the underlying game, and then present the stochastically stable outcome of the imitation process, and comment on its qualitative features. Section 4.4 discusses some other aspects of the results. Section 4.5 concludes.

4.2 Model

Firms are equidistantly located on a Salop circle of unit circumference. Let the number of firms be denoted by n , and we assume that there are at least four of them, i.e. $n \geq 4$. The distance between a firm and its nearest neighbour (on either side) is equal to $\frac{1}{n}$. Firms simultaneously choose price and quantity.⁶

⁵On the other hand, Alós-Ferrer et al. (2000) study an imitation by firms in an evolutionary Bertrand model in a homogeneous market. Firms, which have increasing marginal cost, choose prices and the lowest price-setters win the market, with the market demand function dictating the total quantity to be produced. In this case, the stochastically stable outcome is given by a strict subset of the set of Nash equilibria.

⁶We assume a discrete strategy space by supposing the existence of a price grid and a quantity grid. In particular, we assume that all prominent prices and quantities (that is, those stated in the propositions

Consumers are uniformly distributed on the circumference of the circle. The common intrinsic utility attributed to a unit of the good is denoted by β . They observe the prices posted by the firms, based on which they decide to either purchase the good or to abstain from it. Each consumer acquires at most one unit of the good. On making a purchase, the consumer incurs a linear transportation cost of τ per unit distance to the firm purchased from. The distance between a consumer and a firm or between two firms always refers to the closest distance; due to the circular structure, there exists a clockwise distance and a counter-clockwise distance, and whenever we refer to the distance, we imply the lesser of the two. The net utility of a purchase is the gross utility, β , less the price and transportation cost to the chosen firm. The reservation utility on abstinence from purchase is normalised to 0. A typical consumer maximises net utility. Hence, they buy from a firm only if the net utility from purchase is at least 0; in the case they make a purchase, they buy from the firm that would leave them with the highest net utility.⁷

We assume that, for a consumer, the price chosen by a firm is more salient and observable than the quantity chosen by it, and so the decision of a consumer is based entirely on the price. It is possible that the consumer is not able to acquire the good from the chosen firm due to insufficient quantity produced. We abstract away from the issue of spillover-demand, i.e. we ignore the possibility of a consumer, on not being able to make a purchase from the chosen firm, traveling to another firm that has excess supply. This, in principle, might lead to a situation where one firm faces excess demand while another firm has excess supply. An analysis that incorporates the feature of spillover-demand would necessarily have to come with additional assumptions on how a firm facing excess demand rations or distributes the goods amongst consumers who line-up in front of it. To elaborate briefly, it is possible that a consumer prefers acquiring a good at more than one firm to not acquiring one at all while there might be another consumer who would prefer to acquire the good only from a specific firm. The possibility of spillover demand arises only if both these consumers prefer the same quantity-constrained firm and the second consumer gets served. To examine this, there needs to be an assumption on which consumer acquires the good in case their preferred firm faces excess demand. Even though it might be interesting to study this possibility, we do not pursue it in this chapter.⁸

and needed to complete the arguments) are contained in these grids. Firms can choose any price and quantity from the respective grid.

⁷When the maximum net utility that a consumer receives on purchase is exactly equal to 0, we assume the good is purchased. If a consumer receives the maximum net utility (of at least zero) from more than one firm, we assume that the consumer chooses randomly amongst the firms in contention and there is positive probability of purchase from each of these firms.

⁸However, it is important to mention that our results are robust to one commonly used rationing mechanism in situations of excess demand: efficient rationing. Under efficient rationing, consumers get served in order of decreasing net utility. The reader is also referred to Section 4.4 for a slightly more

The demand (d_i) faced by firm i may be different from the quantity (q_i) produced by it. The amount sold by the firm (s_i) is defined to be the minimum of q_i and d_i . We assume that each firm has a constant cost of production of c per unit and that $\beta > c$ such that there is scope for efficient trade. A firm knows its own marginal cost and does not consider choosing prices below it. The profit of firm i when it chooses a price of p_i and produces q_i is:

$$\pi_i = p_i \cdot s_i - c \cdot q_i.$$

The demand that a firm faces depends on the price that it chooses and the price chosen by the other firms; it comes from the set of consumers who receive the highest net utility from it, subject to the net utility being at least equal to 0. First of all, if the price that firm i chooses is higher than the gross utility, β , then it faces zero demand as any consumer buying from it would receive a net utility that is lower than the reservation utility of not buying. Secondly, firm i also faces zero demand if the closest firm (and so a distance of $\frac{1}{n}$ away from it) chooses a price less than $p_i - \frac{\tau}{n}$; and more generally, it faces zero demand if a firm that is at a distance of $\frac{k}{n}$ away from it chooses a price below $p_i - k\frac{\tau}{n}$. In this case, a consumer never gets the highest net utility from firm i ; given the prices, even the consumer closest to firm i would prefer to incur a higher transportation cost and buy from some other firm.⁹

Having outlined the conditions under which a firm faces zero demand, we now discuss the conditions under which each firm faces positive demand. This implies that all prices are less than β , and are such that a firm at a distance of $\frac{k}{n}$ from firm i chooses a price above $p_i - k\frac{\tau}{n}$, and this holds true for all relevant values of k and $i = 0, 1, \dots, n - 1$. We define a “marginal consumer” to be a consumer who either (a) receives a maximum net utility exactly equal to 0, or (b) receives the same maximum non-negative utility from more than one firm, and so is perfectly indifferent amongst these firms. Two or more firms are said to “compete for the marginal consumer” if they set prices such that there exists a consumer that receives the same maximum non-negative utility from these firms and so, is perfectly indifferent amongst them. If firms do not compete for the marginal consumer, they are said to have segregated markets. We focus first on the demand of a firm when all firms adopt the same price, and then discuss the demand of a firm when all other firms choose the same price but it chooses a different price.

When all firms set the same price p , then the demand faced by each firm is: (i) equal to $2\frac{\beta-p}{\tau}$ if $p > \beta - \frac{\tau}{2n}$ (in this case the firms have segregated markets) or (ii) equal to $\frac{1}{n}$ if $p \leq \beta - \frac{\tau}{2n}$ (here, firms compete for the marginal consumer). On the other hand, if an

detailed discussion on this.

⁹This property would not hold with, for example, firm specific transportation cost.

individual firm chooses price p' when all other firms choose price p , then the demand that the former faces is: (i) equal to $2\frac{\beta-p}{\tau}$ if $p' > 2\beta - p - \frac{\tau}{n}$ (in this case the individual firm has a segregated market) or (ii) equal to $\frac{p-p'+\frac{\tau}{n}}{\tau}$ if $p' \leq 2\beta - p - \frac{\tau}{n}$ (here, firms compete for the marginal consumer).

The profit that a firm receives is determined by the price it chooses and the amount it is able to sell, where the latter (as discussed earlier) depends on the price it chooses, the price other firms choose and the quantity it produces. Since we assume that the consumer decision is based on prices alone and we ignore spill-over demand, the quantity produced by the other firms does not affect its profit. In this chapter, we assume information constraints on a firm. It is cognizant of its own marginal cost and is not aware of other profit-relevant conditions such as the size of the market, the consumer valuation for the product, the preferences of the consumers for different firms (or in other words, the transportation cost), the location of the other firms and the degree of differentiation of the market. The strategy-choice of each firm and the profit it receives is observable. This leads a firm to imitate the firm that receives the highest profit and this situation ensues recurrently. The next section describes the outcome of the imitative process.

4.3 Results

Our focus is on the long-run outcome that emerges from this process of imitation. However, for the purpose of comparison, we begin by presenting the symmetric pure Nash equilibrium of this model and then proceed to the results that directly pertain to the evolution of imitation behaviour.

4.3.1 Nash equilibrium

Even though firms choose price and quantity, for the purpose of the Nash equilibrium, it is sufficient to concentrate on the situation where the firms choose prices. This is because excess supply or demand cannot be a feature of a pure Nash equilibrium. After all, given the strategies of the other firms, a firm can unilaterally increase profit by correcting for excess demand or excess supply.¹⁰ The proposition that follows describes the pure symmetric Nash equilibrium (the proof is in the appendix), which is in fact a strict Nash equilibrium.

¹⁰To be more precise, this would be true if the Nash equilibrium involves price(s) above marginal cost. If, on the contrary, the Nash equilibrium were characterised by all firms choosing marginal cost, then excess demand might be a feature of a Nash equilibrium because correcting for excess demand does not lead to strictly higher profit. As we will see in the proposition to follow, the Nash equilibrium does not involve marginal-cost pricing, and so, the assertion made in the text is true.

Proposition 4.1. (i) Suppose $\tau > n(\beta - c)$. The unique symmetric pure strategy Nash equilibrium is described by all firms choosing the monopoly price and quantity $(p^m, q^m) = (\frac{\beta+c}{2}, \frac{\beta-c}{\tau})$.

(ii) Suppose $\tau \in [\frac{2}{3}n(\beta - c), n(\beta - c)]$. The unique symmetric pure Nash equilibrium is described by all firms choosing the price and quantity $(p^n, q^n) = (\beta - \frac{\tau}{2n}, \frac{1}{n})$.

(iii) Suppose $\tau < \frac{2}{3}n(\beta - c)$. The unique symmetric pure Nash equilibrium is described by all firms choosing the price and quantity $(p^c, q^c) = (c + \frac{\tau}{n}, \frac{1}{n})$.

The proposition reveals that the Nash equilibrium price varies with the magnitude of the transportation cost τ . In addition, when $\tau = n(\beta - c)$, then $p^m = p^n$ (i.e. $\frac{\beta+c}{2} = \beta - \frac{\tau}{2n}$); and when $\tau = \frac{2}{3}n(\beta - c)$, then $p^n = p^c$ (i.e. $\beta - \frac{\tau}{2n} = c + \frac{\tau}{n}$). This implies that the equilibrium price is a continuous function of τ . Further, the equilibrium price increases with the transaction cost up to a value of $\tau = \frac{2}{3}n(\beta - c)$, then up to $\tau = n(\beta - c)$ it decreases, from whereon it does not change with respect to transportation cost. We postpone a discussion of other qualitative features of the Nash equilibrium till Subsection 4.3.3.

4.3.2 Imitation, evolution and stochastic stability

Firms choose price and quantity simultaneously in a situation of recurrent interaction. Since a firm knows its own marginal cost, we assume that it does not price below it. The prices chosen determine the demand that each firm faces while the amount that a firm actually sells is the minimum of the quantity it chooses and the demand it faces, and accordingly, it receives its profit. The price, quantity and profit of each firm is observable, and each firm imitates the actions of the firm that receives the highest profit. If more than one firm receives the highest profit with disparate actions, we assume that there is a positive probability of each of these firms (which receive the highest profit) imitating the other.¹¹ We call this the unperturbed process of imitation.

The state of the process is described by the price and quantity of each firm at each point in time; suppressing the time-index, a state is therefore denoted by $(p_0, q_0, \dots, p_{n-1}, q_{n-1})$. A state is said to be absorbing when there is no possibility of transiting to another state by imitation. A state where all firms choose identical prices and quantities is said to be monomorphic. The specified imitation dynamic gives rise to the following proposition.

Proposition 4.2. *The set of absorbing states consists precisely of all monomorphic states.*

Proof. A monomorphic state is clearly absorbing as firms receive equal profits and have the same actions – this leaves no further scope for imitation. On the other hand, if we are

¹¹This assumption is only to simplify what follows. The set of absorbing states (to be defined soon) would be larger if we assume firms which receive the highest profit (possibly with different strategies) hold on to their strategy, but it has no bearing on the stochastically stable outcome.

in a state where firms realize unequal profits, then firms that realise lower profit imitate the firms that receive the highest profit; if we are in a state where firms realize equal profit with disparate actions, the positive probability of imitation makes this state transient. Thus, from a non-monomorphic state, the process converges to a monomorphic state in finite time.¹²

We now introduce the possibility of firms experimenting with their strategies. Each firm, with an independent probability ε , chooses a random price-quantity pair. This is meant to represent a strategy revision rule that is not necessarily guided by imitation of the most successful firm. Without experimentation, the unperturbed imitation process would converge to (and would stay locked into) an absorbing state that depends on the initial state. Experimentation makes it possible to transit from an absorbing state of the unperturbed imitation process. For example, if one firm experiments with a particular price-quantity choice and, on doing so, receives a higher profit than all other firms, then the other firms imitate the action of the experimenting firm. However, it is possible that the transition from an absorbing state of the unperturbed process requires more than one experimenting firm. Due to the possibility of experimentation, the experimentation-augmented imitation dynamic (or the perturbed process) has no absorbing states. Henceforth, whenever we mention absorbing states, we refer to the absorbing states of the unperturbed process. Stochastic stability is determined by the relative ease of transiting in and transiting out of absorbing states.¹³

The next proposition reveals that the strategies in the stochastically stable state of the imitation dynamic coincide with the Nash equilibrium strategies. The proof of the proposition also explains the reason behind the equivalence of Nash equilibrium and the stochastically stable state.

Proposition 4.3. *In the stochastically stable state, all firms choose the same prices and quantities as in the symmetric pure Nash equilibrium.*

Proof. Let us represent the stochastically stable state by (p^*, q^*) where (p^*, q^*) depends on the transportation cost. First, in (a), we show that it takes a series of single-firm experimentations to get to the state (p^*, q^*) . Next, in (b), we show that it requires more than a single-firm experimentation to transit out (p^*, q^*) .

Part (a). Suppose we are in an absorbing state where each firm chooses $(p, q) \neq (p^*, q^*)$ and obtains profit equal to π . We focus primarily on $p \neq p^*$.¹⁴ Then, we will see that if a

¹²In particular, a monomorphic state where all firms receive identical negative profits is absorbing as well because imitation does not lead to a transition to another state.

¹³For a more detailed exposition of stochastic stability, see, for example, Young (1993) or Kandori et al. (1993).

¹⁴The situation where $p = p^*$ but $q \neq q^*$ is easily addressed. If a firm experiments with q^* , keeping price

firm experiments with p^* and quantity equal to the demand it faces, it obtains a strictly higher profit $\pi' > \pi$ (by the previous proposition). If $p^* > p$, then the profit of the other firms is not affected, and they continue to receive profit π . Note that the neighbouring firms face a higher demand due to a higher price experimentation, but as the quantity is produced in-advance, they cannot take advantage of this. On the other hand, if $p^* < p$, then the experimentation with p^* either affects the other firms negatively by drawing away their consumers and leaving them with excess supply, or does not affect the other firms at all. In any case, the highest profit a non-experimenting firm may receive is π . Since $\pi' > \pi$, the experimenting firm is the most profitable firm, and the price p^* along with the quantity chosen by it is imitated.¹⁵ Subsequently, if a firm experiments with the quantity q^* , which is the actual demand that each individual firm faces given the prevailing price p^* , it obtains the highest profit and is imitated by all the other firms. This brings the process to the absorbing state where each firm chooses the pair (p^*, q^*) ; this happens with a series of single firm experimentations.

Part (b). Suppose we are in the absorbing state where each firm chooses (p^*, q^*) , and all firms receive profit equal to π^* . Proposition 4.1(i) shows that if a firm deviates (in this context, experiments) with any price and quantity, then it constitutes a profit-reducing deviation for that firm. Let the experimenting firm receive $\pi < \pi^*$. First note if a firm experiments with a higher price, it does not affect the profit of the other firms; then, it follows from $\pi < \pi^*$ that the experimentation is not successful. Next, let us take up experimentations with lower price. Even if a firm experiments with marginal cost, it cannot attract any consumers from a firm that is not an immediate neighbour.¹⁶ As a result, the profit of the non-neighbouring firms – the existence of which is guaranteed by the assumption of at least four firms – is not affected by the experimentation, and such firms continue to receive π^* . In either case, the experimenting firm obtains a lower profit than a non-experimenting firm. So, the experimentation is not successful. ■

In part (b) of the proof above, we have mentioned that the existence of at least four firms in protecting against the spite of lower price experimentations when all firms choose the Nash equilibrium price. We will now show that if instead, the market comprised of only

unchanged, it improves its profit as it produces the market-clearing quantity when the price is higher than marginal cost. This makes it the most profitable firm, and it is imitated by all.

¹⁵We show in the proof of Proposition 4.1 that the firm on deviating to p^* in combination with the market-clearing quantity, receives a higher profit than all other firms. This implies that the firm has some freedom in experimenting with the quantity; it does not strictly require to be the market-clearing quantity.

¹⁶This can be verified by the fact that when all other firms choose p^* and a firm experiments with the lowest feasible price of marginal cost, then it can at best provide non-zero utility to the consumer located at the position of the neighbouring firms, i.e. $\beta - p^* - \frac{\tau}{n} \leq 0$. As the non-neighbouring firms do not cater to the consumer located at the position of the neighbouring firm of the experimenter, the former is not affected.

three firms, then an experimentation with a lower price can result in a transition from the Nash equilibrium state. To see this, suppose $\tau < \frac{2}{3}n(\beta - c)$ and all firms choose $p^c = c + \frac{\tau}{n}$. If a firm experiments with $p' = c + \frac{\tau}{2n}$, then the consumers served by either of the non-neighbouring firms equals $(\frac{p'-p^c+\tau/n}{2\tau} + \frac{1}{2n})$. This is because the non-experimenting firms, which produce $\frac{1}{n}$ units, split the consumer located between them, and so, serve $\frac{1}{2n}$ measure of consumers in this part of the market. On the other hand, in the market segment that between the experimenting firm and a non-experimenting firm, the latter serves $\frac{p'-p^c+\tau/n}{2\tau}$ measure of consumers. Thus, the non-experimenting firms are left with excess supply, i.e. $\frac{1}{n} > \frac{p'-p^c+\tau/n}{2\tau} + \frac{1}{2n}$. The profit of the experimenting firm relative to the non-experimenting firm equals $(p' - c)(\frac{p^c-p'+\tau/n}{\tau}) - (p^c(\frac{p'-p^c+\tau/n}{2\tau} + \frac{1}{2n}) - c\frac{1}{n}) > (p' - c)(p^c - \frac{p'+\tau/n}{\tau}) - (p^c - c)(\frac{p'-p^c+\tau/n}{2\tau} + \frac{1}{2n}) = 0$, when $p' = c + \frac{\tau}{2n}$. Thus, the experimenting firm is at least as profitable and is imitated by the other firms. The fact that a single experimentation leads to a transition from the Nash equilibrium state leads to the conclusion that it is not the unique stochastically stable state.

We close this sub-section with a comment on the relationship between the stochastically stable Nash equilibrium and finite evolutionarily stable strategy (FESS). Schaffer (1988) generalises the ESS for the finite population case in the following manner. A population strategy is FESS if the payoff of population strategy playing non-experimenting firms is higher than that of an experimenting firm that chooses a different strategy. Schaffer (1988) goes on to show that FESS and Nash equilibria may not coincide, the reason being that relative payoffs as opposed to absolute payoffs are material for FESS. As a result, strategies that are FESS have to be immune to spiteful actions by mutants – an aspect that is not relevant for Nash equilibrium. However, interestingly, under certain conditions, Alós-Ferrer and Ania (2005) and Schipper (2003) find a close connection between FESS and stochastically stable states under an imitate-the-best dynamic.¹⁷ Further, Tanaka (2000, 20001) shows that in a differentiated market (with linear demand functions), the stochastically stable state with price or quantity imitation is FESS.

The concept of FESS, as defined by Schaffer (1988), is not very applicable to the situation at hand because the profits (or payoffs) depend not only on the strategy profile, but also on the location, which is not a strategy here. In presence of an experimenting firm, the average profit of the non-experimenters may be different from the minimum and maximum profit of the non-experimenters. This may suggest for three different notions of FESS where the profit of the mutant is compared to these three statistics i.e. the average,

¹⁷Anonymity and “aggregate-taking”, i.e. the payoff of an economic unit depends only on the strategy of that unit and the aggregate strategy of the other units, play an important role in both contribution. In contrast, anonymity is not a property of our situation: due to the spatial structure, the strategies of the other firms are not sufficient to determine the payoff of a firm; the payoff of a firm, in general, also depends on the identity of the firms.

minimum or maximum profit. While in general, the payoff of the mutant may be higher than the average profit (and hence the minimum profit) of the non-mutants, it is not higher than the maximum profit of the non-mutants. Thus comparison of the profit of the mutant with the average or minimum profit of non-mutants would reveal the Nash equilibrium or stochastically stable states to not be FESS, while comparison with the maximum profit of non-mutants would reveal it to be FESS.

4.3.3 Qualitative features of the (long-run) equilibrium

The equilibrium price is a continuous function of transportation costs τ . It increases with this cost up to a value of $\tau = \frac{2}{3}n(\beta - c)$ to a price level above monopoly prices, from where it decreases to reach the monopoly price at $\tau = n(\beta - c)$, and settles there for higher values of the transportation costs (see Figure 4.1).

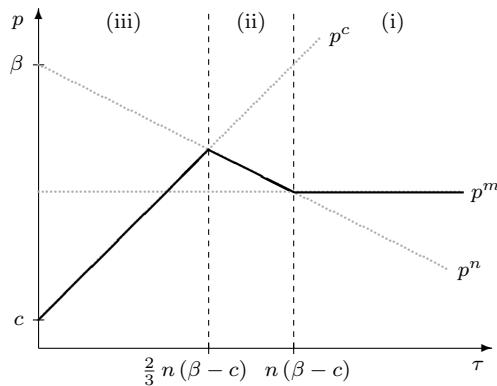


Figure 4.1: Stochastically stable state prices as a function of transportation costs.

When transportation costs are high as in area (i), the firms charge the price that they would charge if they were a monopolist (on the Salop circle). The market is not fully served in the sense that there is a positive mass of consumers who prefer not to acquire the good. This implies, in turn, that firms do not compete for the marginal consumer and this is driven by the fact that the transportation costs are too high. High transportation costs make it difficult for the firms to attract consumers not close enough with a price at least equal to marginal cost. This gives each firm some degree of market power, which is manifested in monopoly prices and quantities in the stochastically stable state. If firms were to try to serve the entire market, then the price would have to be reduced substantially (possibly even below marginal cost level), and the increase in the market size would not justify (in terms of relative or absolute profit) this reduction in price.

When transportation costs decrease to the range in area (ii), the firms choose a price such that the entire market is served (i.e. all consumers acquire the product). They price “non-aggressively” in the sense that they share the market while extracting all surplus

from the marginal consumers, i.e. the consumers located exactly in between two firms (at a distance of $\frac{1}{2n}$ from either firm) receive a net utility of zero. While this may seem collusive, with prices being above monopoly prices, it actually is not: A unilateral price decrease, along with the optimal quantity adjustment, by one firm is (absolutely and relatively) profit-reducing.¹⁸ So, a decrease in price (which would lead to ‘real’ competition amongst firms) neither leads to a higher absolute profit as in the Nash-equilibrium paradigm, nor to a (weakly) higher profit relative to the other firms as in the context of the imitation paradigm. Within this range of transportation costs, the price is actually decreasing with the transportation costs, implying that more market differentiation may lead to lower prices. This striking phenomenon is caused by the firms extracting all excess surplus from the marginal consumers (whose utility is decreasing in the transportation cost) in a fully served market.

As transportation costs drop beyond the threshold dividing area (ii) and (iii), the equilibrium price is given by a mark-up over marginal cost. The market is fully-served and all consumers enjoy a positive net surplus. For transportation costs above $\frac{1}{2}n(\beta - c)$, the price is above monopoly price¹⁹; for lower values the price is below monopoly price level. As the market becomes homogeneous (i.e. τ approaches zero), the price approaches the competitive price.

The effect of an increase in the number of firms (n) on the equilibrium price is depicted in Figure 4.2. The solid (dotted) line depicts prices as a function of the transportation costs for smaller (larger) number of firms. For high levels of the transportation costs, the firms charge the price that they would charge if they were a monopolist and, hence, the price does not depend on the number of firms. Interestingly, for middle ranged values of the transportation costs, the dotted line is above the solid line, meaning that, *ceteris paribus*, prices increase with the number of firms. This is a reflection of the fact that for this range the market is divided among firms and maximal surplus is extracted from consumers. As the number of firms increase, the marginal consumers (i.e. the consumers located exactly in between two firms) are closer to the firms on either side, and in order to fully extract these consumers’ surplus, the price has to be higher. Thus, more competition may actually lead to higher prices.²⁰ For low levels of the transportation costs, we recover the standard

¹⁸For literature on the attainability of collusion in differentiated markets, see Colombo (2012).

¹⁹For (iii), $c + \frac{\tau}{n} > \frac{\beta+c}{2}$ if and only if $\tau > \frac{1}{2}n(\beta - c)$.

²⁰Existing literature on price increasing competition include Satterthwaite (1979) and Stiglitz (1987), where the presence of search costs results in a higher number of sellers leading to higher prices; Rosenthal (1980) makes the same observation when there exists two sets of consumers – one set of consumers who are loyal to a seller while another group of consumers switch between sellers. Our set-up does not involve search costs or groups of consumers, and so, our results are closer to Chen and Riordan (2008), who analyse general conditions when product differentiation leads to prices under which a duopoly being higher than price under a monopoly.

intuition that more competition leads to a lower price. The range of transportation costs for which this subcase is valid is increasing in the number of firms. In fact, when the number of firms grows large, prices converge pointwise to marginal cost (for all levels of transportation costs).

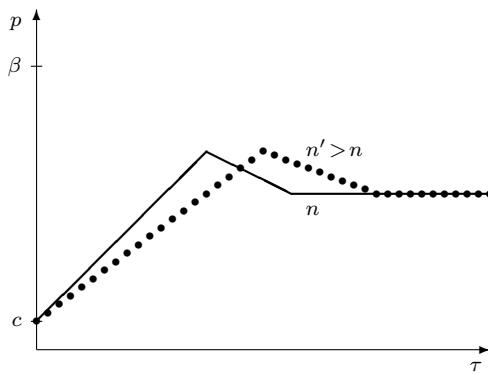


Figure 4.2: Relation between number of firms and stochastically stable state prices.

4.4 Discussion

We begin with a comparison of the presented results with the long-run outcome of imitation driven behaviour when firms choose prices but production is on-demand. First, we stress on a crucial difference between production on-demand and production in-advance. With production in-advance, when an experimenting firm chooses a lower price, the neighbouring firms see some of their consumers taken away, and so, are left with excess supply. This implies that the non-neighbouring firms of the experimenting firm are the most successful non-experimenters. On the other hand, production on-demand would exclude excess supply by assumption. However, as far as the profit comparison between an experimenting firm and the most successful non-experimenting firm is concerned, there is no difference between these two approaches as the most successful non-experimenting firm is one that is not affected by the experimentation. But, when an experimenting firm chooses a higher price, the neighbouring firm would face excess demand when production is in-advance; in this case, all non-experimenting firms have the same profit as they sell the same amount at the same price. Production on-demand does not allow for excess demand, and now, the profit of the neighbouring (non-experimenting) firms of the experimenting firm may be more profitable than other non-experimenting firms. Thus, comparing the relative profit of the experimenting firm and the non-experimenting firm may yield different conclusions depending on whether production is in-advance or on-demand.

We are now in a position to discuss this difference; we only discuss the intuition for

the results – the exact working of imitation with production on-demand is slightly more involved. With production on-demand, for high transportation costs, the stochastically stable state coincides with Nash equilibrium behaviour. The reason for this is similar to the one in subsection 4.3.3: high transportation cost results in Nash behaviour being stable in an imitation paradigm. For intermediate values of transportation cost, Nash behaviour is still stochastically stable state but not uniquely so; monomorphic states where all firms post the same price from a particular set of prices below the Nash equilibrium, but higher than marginal cost, are also stochastically stable. The reason for this is two-fold: firstly, moderate transportation cost results in the spite effect of lower prices having more bite, and secondly, the Nash equilibrium price is not competitive enough to withstand this bite. However, for low values of transportation cost, Nash behaviour is again uniquely stochastically stable (when there are at least seven firms). The reason for re-obtaining Nash behaviour as the unique stochastically stable state is simply that the Nash equilibrium price is competitive enough to withstand the spiteful effects of lower price experimentations, even though lower transportation costs create more conducive conditions for spiteful behaviour. In summary, we see that Nash behaviour is always stochastically stable; further, except for a range of intermediate values of transportation costs, the long-run behaviour coincides uniquely with the Nash equilibrium behaviour, and this is independent of production being in-advance or on-demand.

This leads us to a related point, i.e. the manner in which we treat the situation of excess demand with production in-advance. While we have assumed away the possibility of spill-over demand, this assumption does not drive the results. One commonly rationing mechanism in case of excess demand is efficient rationing, where consumers are served in decreasing order of net utilities till the point that demand can no longer be met. The results presented here are robust to the possibility of efficient rationing. This is because excess demand does not influence (i.e. cause or prevent) in any way the transition from absorbing states that are not stochastically stable to the stochastically stable state, or vice-versa. Hence, ignoring spill-over demand is not vital to our analysis.

Apesteguia and Selten (2005) provide experimental evidence of imitation behaviour when three, four or five firms on a Salop circle set prices, and receive feedback on the performance of the immediate neighbours. Their interpretation is that in spite of behaviour being consistent with imitation, the outcome does not exactly correspond to the Nash equilibrium, and they develop the concept of an “imitation equilibrium” to explain the observed outcome. Another plausible explanation for the suggested disparity of the observed outcome from the Nash equilibrium may be that in the experiment each firm observes only the two immediate neighbours (as a result of which, a firm may not be able to imitate the most successful firm but only the most successful firm in the neighbourhood). It is possible

that imitation under incomplete observation, where a firm only imitates the most successful firm it can observe, leads to a different outcome from imitation under complete observability.²¹ Further, in some situations there are only three firms, and we have mentioned the importance of at least four firms in sustaining our result.

The results of this chapter complement Hohenkamp and Wambach (2010), who study the evolution of firm location by imitation. In their paper, the evolution of firm location is modeled as a first stage game, with the price setting by firms being the second stage game. In the second stage where firms set prices, they assume firms to exhibit Nash equilibrium behaviour. It is then argued that evolution would lead to firms being located equidistantly on the Salop circle. Here, we fix the location of the firms and analyse the evolution of price and quantity, and show that the emergence of Nash equilibrium in the price–quantity space. In combination with Hohenkamp and Wambach (2010), this hints at the possibility that evolution by imitation leads to Nash equilibrium outcomes in terms of location, as well as in terms of price and quantity. Thus, a simple imitation rule that also finds some experimental support (see for example, Huck et al. (1999)), exhibits Nash equilibrium behaviour in both location and price.

We conclude with a comment on incorporating heterogeneity in firm behaviour. What would be the effect of including, in addition to imitating firms (imitators), optimising firms that choose prices and quantities that maximise their profit given the strategies of the other firms? There are at least two ways of incorporating such heterogeneity, and we briefly comment on each possibility in turn. Firstly, assume that each firm, with positive probability, either imitates the firm with the highest profit, or plays a best response to the strategies chosen by all other firms. The motivation is that, in each period, a firm may appoint a decision-maker who, with some probability, imitates or optimises. The stochastically stable state is robust to this extension. The simplest way of arguing for it is by observing that now, the symmetric Nash equilibrium state is the only absorbing state. This is because with positive probability, each firm exhibits only imitative behaviour, and so, only monomorphic states may potentially be absorbing. But, if a firm were to best-respond to the other firms' strategies, then the only state that can be absorbing with respect to imitation and myopic best-response would be the unique symmetric Nash equilibrium; thus, a myopic best response will drive the system away from any other monomorphic state. Hence, there exists only one absorbing state: the unique symmetric Nash equilibrium; and so, it is (trivially) the only stochastically stable state. In this state, optimisers and imitators earn the same profit, and there is no evolutionary pressure that acts against either of the two behavioural rules.

²¹Boncinelli (2008) and Alós-Ferrer and Weidenholzer (2010) find that with imitation behaviour in coordination games, incomplete observation may have an effect on the stochastically stable outcome.

Another way of incorporating this two-type behavioural heterogeneity is by specifying that an individual firm either always imitates, or always optimises. The difference between this conceptualisation and the one in the preceding paragraph is that here, there is no within firm behavioural heterogeneity; rather, the behavioural heterogeneity is across firms. This is similar in spirit to Hohenkamp and Kaarbøe (2008), Josephson (2009) and Schipper (2009), who analyse the long-run outcome when a population contains both imitators and optimisers. The analysis of this situation is more involved, for the reason that the strategy chosen by an optimiser depends not only on the strategies chosen by the other firms, but also on the location and behavioural type of the other firms. For example, the strategy of an optimiser with imitators as its immediate neighbours will, in general, be different from the strategy of an optimiser that has one optimiser and one imitator as its neighbours, and this, in turn, may be different from the strategy of an optimiser that has two optimisers as neighbours. As a result, the spatial arrangement of the behavioural types is a material consideration, and adds to the complexity of the analysis.

4.5 Conclusion

We present an evolutionary model of firm behaviour in a horizontally differentiated market. Firms are equidistantly located on the circumference of a unit circle, while consumers are distributed uniformly along it. A firm chooses price and quantity, while consumers make their decision to purchase based on the announced prices. We assume a very limited informational environment: a firm is assumed to know only its own marginal cost. At the end of each period, the strategy and profit of each firm is observable and a firm imitates the strategy of the most successful firm, while occasionally experimenting with random strategies.

The stochastically stable outcome of this process depends on the level of differentiation and coincides with the Nash equilibrium of the underlying game, in spite of the simple heuristics used by a firm. When the level of differentiation is high, in the stochastically stable outcome, firms choose monopoly prices and outputs. At moderate levels of differentiation, firms price such that the surplus of the marginal consumers is fully extracted; nevertheless, the entire market is served. At low level of differentiation, firms price at a mark-up over marginal cost and the entire market is served. In all cases, market-clearing arises endogenously. Interestingly, for some range of differentiation levels, an increase in the number of firms may lead to higher prices and an increased level of differentiation may lead to lower prices.

4.6 Appendix

Proof of Proposition 4.1

Statement (i): $\tau > n(\beta - c)$. First, in (a), we show that each firm choosing the strategy $(p^m, q^m) = (\frac{\beta+c}{2}, \frac{\beta-c}{\tau})$ is a strict Nash equilibrium, by showing that all unilateral deviations are profit-reducing. Next, in (b), we show that there does not exist any other symmetric pure Nash equilibrium.

Part (a). Suppose each firm chooses the monopoly price and quantity pair $(p^m, q^m) = (\frac{\beta+c}{2}, \frac{\beta-c}{\tau})$. Then, firms do not compete for the marginal consumer. Firms would compete for the marginal consumer (while setting the same price) if the prices were below $\beta - \frac{\tau}{2n}$. The restriction $\tau \geq n(\beta - c)$ ensures that $p^m = \frac{\beta+c}{2} \geq \beta - \frac{\tau}{2n}$. So, each firm receives the monopoly profit of $(p^m - c)q^m = \frac{(\beta-c)^2}{2\tau}$.

An important price level when considering a unilateral deviation is $\hat{p} = 2\beta - p^m - \frac{\tau}{n}$. Prices below this level induce competition for the marginal consumer with the neighbouring firm. As it is clear that a unilateral deviation can never increase the profit for a firm deviating to a price above this level, consider a deviation to a price p' below \hat{p} . This deviation induces competition for the marginal consumer with the neighbouring firms and the demand that it faces equals $\frac{p^m-p'+\frac{\tau}{n}}{\tau}$. The deviation would be at least as profitable for the firm if

$$(p' - c) \frac{p^m - p' + \frac{\tau}{n}}{\tau} \geq \frac{(\beta-c)^2}{2\tau}.$$

The left hand-side of the expression (i.e. the profit after deviation) is maximised at $p^* = \frac{1}{2} \cdot (p^m + c + \frac{\tau}{n})$. However, $\tau \geq n(\beta - c)$ implies $p^* > \hat{p}$, and so the price p^* cannot induce competition for the marginal consumer. So, we have to consider the deviating price that maximises the profit on the left-hand side while satisfying this assumption. From the shape of the profit function of the firm on deviation to price p' , it follows that the constrained profit-maximising price is given by $p^\circ = 2\beta - p^m - \frac{\tau}{n}$. It can be easily verified that for $p' = p^\circ$ the inequality above cannot be satisfied. Hence, even the best possible deviation decreases a firm's profit, leading to the conclusion that the state where all firms choose (p^m, q^m) is a strict Nash equilibrium.

Part (b). Suppose that we are in a state where all firms choose the same price and quantity (p, q) . First, suppose $p > 2\beta - p^m - \frac{\tau}{n}$ and $p \neq p^m$. If a firm deviates to (p^m, q^m) , it does not compete for the marginal consumer and earns monopoly profit, which must be higher than the pre-deviation profit. Hence, monomorphic states with $p > 2\beta - p^m - \frac{\tau}{n}$ and $p \neq p^m$ cannot be Nash equilibria.

Now, consider $p \leq 2\beta - p^m - \frac{\tau}{n}$. The assumption $\tau \geq n(\beta - c)$ implies $p \leq 2\beta - p^m - \frac{\tau}{n} \leq \beta - \frac{\tau}{2n}$. Consequently, at prices p , firms compete for the marginal consumer. The demand

that they face equals $\frac{1}{n}$ and they receive a profit of $\frac{p-c}{n}$. Now, suppose a firm deviates to the monopoly price p^m . As this deviating firm still competes for the marginal consumer, the demand it faces equals $\frac{p-p^m+\frac{\tau}{n}}{\tau}$; let the firm also choose this quantity. The profit from the deviation is higher than the profit before deviation if

$$(p^m - c) \frac{p-p^m+\frac{\tau}{n}}{\tau} - (p - c) \frac{1}{n} > 0.$$

The partial derivative of the left hand-side expression with respect to p is negative, meaning that it is minimised at the highest value of p subject to the constraint $p \leq 2\beta - p^m - \frac{\tau}{n}$. So, if the inequality is satisfied for $p = 2\beta - p^m - \frac{\tau}{n}$, it is also satisfied for all lower values of p . When $p = 2\beta - p^m - \frac{\tau}{n}$, the inequality can be rewritten as $\frac{(\beta-c)^2}{2\tau} > (p - c) \frac{1}{n}$. As the left hand-side equals monopoly profit – which is the maximum a firm could ever get – and the right hand-side represents a feasible profit, this inequality is always satisfied. We conclude that monomorphic states with $p \leq 2\beta - p^m - \frac{\tau}{n}$ cannot be Nash equilibria.

Statement (ii): $\tau \in [\frac{2}{3}n(\beta - c), n(\beta - c)]$. First, in (a), we show that each firm choosing $(p^n, q^n) = (\beta - \frac{\tau}{2n}, \frac{1}{n})$ is a strict Nash equilibrium. Then, in (b), we show that there does not exist any other symmetric pure Nash equilibrium.

Part (a). Suppose that each firm chooses (p^n, q^n) . The profit that each of them receive equals $\frac{1}{n}(p^n - c)$. If a firm deviates with a price-quantity pair (p', q') with $p' > p^n$, then the deviating firm does not compete for marginal consumers. As a result, the demand the firm faces equals $2\frac{\beta-p'}{\tau}$ and the profit equals $2(p' - c)\frac{\beta-p'}{\tau}$, which is decreasing in p' for $p' \geq \frac{\beta+c}{2}$. From $\tau < n(\beta - c)$, it follows that $p^n = \beta - \frac{\tau}{2n} > \frac{\beta+c}{2}$. So, a deviation to a price $p' > p^n$ implies $p' > \frac{\beta+c}{2}$ – so, the firm receives a strictly lower profit. Thus, there does not exists a profitable deviation for a firm when it chooses a higher price.

So now, let a firm deviate with a different price-quantity pair (p', q') with $p' < p^n$. The deviating firm now steals away some consumers of the two neighbouring firms and faces a demand of $\frac{p^n-p'+\frac{\tau}{n}}{\tau}$; let this also be the quantity the firm chooses. This will represent at least as profitable deviation only if

$$(p' - c) \frac{p^n-p'+\frac{\tau}{n}}{\tau} - (p^n - c) \frac{1}{n} \geq 0.$$

The left hand-side expression is maximised for $p^* = \frac{1}{2}(p^n + c + \frac{\tau}{2n})$. But $\tau > \frac{2}{3}n(\beta - c)$ implies $p^* > p^n$. As the expression above is globally concave and is maximised at $p^* > p^n$, it follows that a deviation to a price lower than p^n results in strictly lower payoff. So, all deviations are profit reducing, and so the state where all firms choose (p^n, q^n) is a strict Nash equilibrium.

Part (b). Suppose we are in the state where all firms choose the same price-quantity pair (p, q) with $p > p^n$. Then, the firms do not compete for the marginal consumer. The

demand and profit each firm faces equals $2\frac{\beta-p}{\tau}$ and $2(p-c)\frac{\beta-p}{\tau}$ respectively. This profit is decreasing in p if $p > \frac{\beta+c}{2} (= p^m)$. Since $p^m < p^n < p$, a firm deviating to (p^n, q^n) will increase its profit. Hence, monomorphic states with $p > p^n$ are not Nash equilibria.

Suppose instead that all firms choose the price-quantity pair (p, q) with $p < p^n$. Now, firms compete for the marginal consumer. Since all firms choose the same price, each firm faces a demand of $\frac{1}{n}$ and receives a profit of $\frac{1}{n}(p - c)$. Consider a firm deviating to (p^n, q') . As this firm chooses a higher price than its neighbouring firms, it loses some of its consumers and will be left with $\frac{p-p^n+\frac{\tau}{n}}{\tau}$ consumers.²² Suppose that the deviating firm adapts the quantity optimally given the price: $q' = \frac{p-p^n+\frac{\tau}{n}}{\tau}$. This would be a profitable deviation if

$$(p^n - c) \frac{p-p^n+\frac{\tau}{n}}{\tau} - (p - c) \frac{1}{n} > 0.$$

The derivative of the left hand-side expression with respect to p is $\frac{1}{\tau}[p^n - (c + \frac{\tau}{n})]$, which is negative as $\tau > \frac{2}{3}n(\beta - c)$. Under the constraint $p < p^n$, the expression is minimised at $p = p^n$. At $p = p^n$, it is exactly equal to 0. Therefore, it is positive for all other prices $p < p^n$. Thus, there exists a profitable deviation from monomorphic states with $p < p^n$, and so, such states are not Nash equilibria.

Statement (iii): $\tau < \frac{2}{3}n(\beta - c)$. See, for example, Tirole (1988). ■

²²The deviating firm always faces a strictly positive demand. There would be no demand if and only if $p^n = \beta - \frac{\tau}{2n} > p + \frac{\tau}{n}$, as then all consumers would receive higher utility by acquiring from its neighbouring firm. This condition is satisfied if and only if $\tau < \frac{2}{3}n(\beta - p)$ and this is not satisfied for the lowest possible price of p , which is $p = c$. Hence, the deviating firm faces a positive demand even when other firms price at marginal cost and, so, it faces positive demand when the other firms choose any price $p < p^n$.

Chapter 5

Local Imitation and Coordination¹

5.1 Introduction

This chapter studies the evolution of behaviour in coordination games when individuals choose their strategy by imitation of the most successful individual they observe. The individual's scope of strategic interaction is global, in the sense that the co-player may be drawn from the entire population. However, due to informational or other constraints, the individual is only locally informed, i.e. he is aware of the experience of only a subset of individuals in the population, namely his neighbours. Individuals imitate the most successful individual in their respective neighbourhoods and this may differ from imitating the most successful individual. We investigate the effect of such a “local” imitation protocol when the interaction is in a wider environment.

The set-up has three principal ingredients. Firstly, the strategic nature of the situation, that is recurrent, is described by a 2×2 coordination game where the Pareto-efficient coordination equilibrium differs from the risk-dominant coordination equilibrium.² There is, therefore, a divergence between the preferred equilibrium and the “safer” equilibrium. An individual is randomly matched with another co-player from a finite population. Random matching carries the import that there are circumstances where an individual does not have prior knowledge of the person he will be interacting with. This situation ensues recurrently and so, an individual has to choose a strategy each period.

Secondly, we assume that an individual does not have complete information about the environment that he is a part of. He is able to observe, acquire and exchange information (i.e. strategy choice and payoff received) from a subset of individuals in the entire population, say his neighbours. This is modeled by placing individuals on a network, where the

¹This paper is based on the chapter “Coordination under Global Random Interaction and Local Imitation”, Meteor Research Memorandum 13/004.

²Later, we generalise to $m \times m$ coordination games where the coordination outcomes may be Pareto-ranked.

links of an individual determine his social network and his set of neighbours. The links are assumed to be undirected (i.e. the flow of information between two linked individuals is bidirectional) and exogenously given.³ Even though individuals may not observe everyone in the population, it is likely that any two individuals in the population are connected either directly or indirectly via other individuals. So, we assume the network to be connected.

Thirdly, we assume that an individual, owing perhaps to reasons of complexity, informational constraints or some form of bounded rationality, imitates the most successful individual (including himself) that he observes in the previous period. Apesteguia et al. (2007) find experimental support for imitative behaviour while Björnerstedt and Weibull (1996) provide theoretical support for such behaviour in a game-theoretical context.⁴ In this model, imitation may be interpreted as a means by which an individual tries to grope for the action that might result in a better outcome; and since an individual by himself is not aware of the state of the entire world, he also relies on the experience of his acquaintances.

We ask the following questions: what is the effect of incomplete observation induced local imitation on the stochastically stable outcome when individuals interact globally? When are societies able to coordinate on the Pareto-efficient equilibrium strategy? That is, what are the informational network structures under which a society is able to converge to the state in which all players coordinate on the Pareto-efficient equilibrium?

Online platforms serve to illustrate a situation where an individual may interact with strangers, i.e. other individuals with whom one is not linked or connected to (and thereby cannot acquire information from). An individual may be randomly matched online with such a stranger in situations where it may be preferable but risky to coordinate on a particular strategy (such as choosing the same file format or application). It is possible that individuals get information about payoffs to each action from a narrower set of people than they may interact with online, and they imitate the most successful experience. Another example involves the evolution of socio-cultural features such as punctuality, which may be modeled as a coordination game (as in Basu and Weibull (2003)). Two randomly selected individuals are supposed to meet at a pre-specified time. Coordinating on punctual behaviour is better for both individuals as opposed to a situation when both of them are

³The interpretation is that the links are given exogenously by history and social factors and is inherited by an individual. Of course, there may be situations wherein the process of link formation or deletion may merit as much attention as the choice of strategy. For example, Goyal and Vega-Redondo (2005) discuss the case where not only do individuals play a coordination game, but also establish unilateral costly links with others; these links determine the set of people with whom the interaction occurs. However, here, the interest is not how the information structure evolves via the network; rather, the network is treated as a modeling device that determines information percolation.

⁴Imitation of successful behaviour also finds mention in social and biological literature such as Dawkins (1976) and Ulmann-Margalit (1978) as a propellant of human intelligence.

impunctual. In case of mis-coordination of actions, the punctual individual incurs a cost (for example, of waiting). Thus, both individuals arriving on time represents the Pareto-efficient outcome while both individuals being tardy represents the risk-dominant outcome. An individual selects a particular course of action and receives a payoff depending on his action and that of his randomly selected partner. Individuals exchange their respective experiences with their neighbours and thereafter, an individual selects the strategy that yielded the highest payoff. In either case, our interest lies in the mode of behaviour that is expected to emerge in the long-run. How does the information structure affect the custom that individuals settle on?

A natural benchmark for the model of imitation under incomplete observability presented here, is imitation under complete observability, i.e. when individuals can observe everyone and hence, imitate the strategy choice of the most successful individual in the entire population. Robson and Vega-Redondo (1996) present a similar model, where in each period, every individual interacts with v other randomly chosen individuals to play a coordination game. The actions and payoffs of individuals are observable to all and the action that yields the highest average payoff is imitated. Here, for tractability, we assume that an individual interacts with only one individual per period, and that individuals imitate the individual with the highest payoff. As the action of a more successful individual is more salient than a strategy that gives the higher average payoff, it may be more plausible to assume that in certain situations, the more successful individual is imitated.⁵

The full observability model is a particular instance of the model of local observability. The latter generates the former when the network is complete, i.e. each individual is linked to every other individual. In Lemma 5.1, we show that in either case, the defined imitation rule leads to all individuals choosing the same strategy. So, the process converges to either the Pareto-efficient equilibrium or the risk-dominant equilibrium. To obtain the stochastically stable outcome, infinitesimally small perturbations are introduced in the imitation dynamic. These perturbations may be interpreted as the propensity of individuals to choose a strategy not necessarily prescribed by the specified imitation rule, owing perhaps to tendencies to experiment or make mistakes. This makes the transition from the state where everyone uses the strategy of the Pareto-efficient equilibrium to the state where everyone adopts the risk-dominant equilibrium strategy, and vice versa, possible. The equilibrium that is more resistant to such perturbations is stochastically stable.

In the full observability model, the Pareto-efficient equilibrium is the stochastically stable state (see Remark 1) for the reason that the risk-dominant equilibrium is relatively

⁵An observation consistent with this assumption involves people who desire to be actors or sportpersons. In light of the very small percentage of people who find success, the average payoff for choosing such an action might not be very high. It seems that the attractiveness of such professions is (at least partially) driven by the salience of people who actually end up being successful.

more susceptible to experimentation by individuals. To elaborate, suppose that the state is given by all individuals playing the risk-dominant strategy. If two individuals experiment with the strategy of the Pareto-efficient equilibrium, and are also matched to each other - a random event with positive probability - they obtain the highest payoff. This is observed by everyone and their strategy is imitated. The Pareto-efficient equilibrium is more resistant as, for populations that contain at least six individuals, the transition to the risk-dominant equilibrium requires at least three experimentations.

With limited observability (i.e. when the connected network is not complete), while the Pareto-efficient equilibrium is always stochastically stable, it may not be uniquely so; the risk-dominant equilibrium may be stochastically stable as well. The reason for the stability of the Pareto-efficient equilibrium is similar. If two individuals experiment to adopt the strategy of the Pareto-efficient equilibrium when all else play the risk-dominant strategy, and if these two individuals are also matched, they obtain the highest payoff. This is observed by their neighbours, who then imitate this strategy. With positive probability, the strategy spreads to the neighbours' neighbours and ultimately, because of connectedness of the network, to the entire population in finite time. The reverse transition takes at least two experimentations. (This establishes the stochastic stability of the Pareto-efficient equilibrium and this is the content of Proposition 5.2.) With an example, we show that the reverse transition may be possible with two experimentations. This possibility arises precisely because of limited observability; the risk-dominant equilibrium strategy may spread by imitation because it is the most successful strategy that individuals observe. Hence, this example illustrates that limited observability does make a difference.

A natural follow-up question concerns conditions under which the Pareto-efficient equilibrium is expected to uniquely prevail and, to this end, we identify two alternative (sufficient) conditions. One is that each individual, independent of the size of the population or the payoff structure of the game, needs to observe (i.e. be linked to) at least four other individuals (Proposition 5.3). This can be interpreted as an absolute criterion – a minimum level of observability that, independent of other factors, suffices to support the unique stability of the Pareto-efficient equilibrium. Proposition 5.4 reveals a second condition that deals with regular networks; the Pareto-efficient equilibrium is uniquely stable when the network is regular (i.e. all individuals are linked to the same number of individuals). Further, we present examples to show that the sufficient conditions mentioned above are not very slack.

Related papers which posit an imitation dynamic in a game-theoretic setting (as opposed to imitation in a decision-theoretic framework as in Ellison and Fudenberg (1993, 1995) or Banerjee and Fudenberg (2004)) include Robson and Vega-Redondo (1996) and Eshel, Samuelson and Shaked (1998). The former assumes that an individual interacts

with randomly chosen individuals from the entire population and imitates the strategy that gives the highest average payoff in the entire population. The latter focusses on local structures: not only does an individual interact only with his neighbours, but imitates the strategy that yields the highest average payoff in his neighbourhood. However, irrespective of the scope being global or local, the interaction and information spheres coincide in these papers. Alós-Ferrer and Weidenholzer (2008) make a distinction between these two by assuming that the information sphere is larger than the interaction sphere: while individuals play the coordination game in their neighbourhood, they hear of the behaviour of a larger set of individuals; individuals imitate the most successful individual they know of and this may be an individual they never interact with themselves (adoption of best-practises in an industry is given as an example). The motivation here is different: individuals might realistically not be able to observe everyone in the population, but the person with whom they interact might be randomly chosen. Given that they cannot observe everyone, they imitate the most successful individual they can observe. This results in a situation where the interaction sphere is larger than the information sphere.

In summary, the objective of this chapter is to examine the effect of paucity of information on an imitation-driven process in the context of coordination games on arbitrary networks, where paucity of information is interpreted as the set of people from whom information may be acquired being a subset of the individuals with whom one may interact with. Existing literature does not consider this aspect and deals almost exclusively with specific network architectures. For example, Ellison (1993) develops a local interaction model where individuals located on a circle play a coordination game with their neighbours according to a best response dynamic, the outcome being that the risk-dominant equilibrium is expected to emerge most of the time. It is difficult for us to draw a comparison with this result as it is based on a completely different adjustment rule, i.e. best response instead of imitation. Eshel et al. (1998), Mengel (2009) and Hedlund (2012) deal with a similar local interaction model, based on imitation, but the underlying game is better described as a Prisoner's dilemma rather than a coordination game, thereby limiting the scope of a direct comparison. On the other hand, Chen, Chow, and Wu (forthcoming) impose imitation behaviour on individuals located on a circle in the context of coordination games. They obtain that, at best, the Pareto-efficient equilibrium strategy players co-exist with the risk-dominant equilibrium strategy players, and in other cases, the risk-dominant equilibrium prevails uniquely. This contrasts directly with our result that the Pareto-efficient equilibrium prevails uniquely when each individual observes the same number of other individuals (i.e. the network is regular). However, because Chen et al. (forthcoming) assume that individuals imitate the strategy in their neighbourhood with the highest average payoff (as opposed to the imitation of the strategy in an individual's neighbourhood

that yields highest payoff – an assumption we make because often, payoffs are more salient than strategies), the difference cannot be attributed only to lack of information. Alós-Ferrer and Weidenholzer (2006) under the same local interaction structure observe that when individual imitate the best performing neighbour, the Pareto-efficient equilibrium is stochastically stable when the interactions are neither too global or too local.⁶

In the afore-mentioned papers, individuals receive information from the entire set of individuals they interact with, and so, from the perspective of our suggested interpretation, information is not incomplete. As referred to earlier, Alós-Ferrer and Weidenholzer (2008) assume that individuals gather information from a wider group than they interact with, and in this situation, the Pareto-efficient equilibrium prevails if the interaction is not too global and the flow of information is fluid. In our case, in contrast to Alós-Ferrer and Weidenholzer (2006, 2008), interaction is global in the sense that interaction is not restricted to one's own neighbourhood, while the flow of information is restricted to one's neighbourhood. The most convenient way of observing the bearing of paucity of information (modeled by restriction of information flow to one's neighbourhood) is via a comparison of Remark 1 and Proposition 5.2: with incomplete observability, the Pareto-efficient equilibrium is no longer the unique stochastically stable state.

The structure of the chapter is as follows. Section 5.2 presents the model and preliminaries. The results are contained in Section 5.3. In Section 5.4, we reason that the results presented may be extended readily to coordination games with a larger strategy space (i.e. $m \times m$ coordination games), and that the import of the results are robust even if the imitation protocol exhibits inertia. Section 5.5 concludes. All proofs are presented in the appendix.

5.2 Model and Preliminaries

We consider a finite population comprising of M individuals and denote the set of all individuals by I . In each period, an individual chooses between strategies P and R , after which he is randomly matched to another individual; there is a strictly positive probability of every possible match and it is possible to incorporate the feature that, based on certain factors, some matches are more probable than others. Throughout the chapter, it is assumed that the population is made up of an even number of individuals (to avoid the situation where someone is left unmatched) and contains at least six individuals; so, $M \geq 6$

⁶Weidenholzer (2010) surveys the literature pertaining to coordination games and local interactions.

and even.⁷ The paired individuals interact in a 2×2 symmetric coordination game.⁸ The payoff that an individual receives on choosing strategy s_i is denoted by $u_i(s_i)$ (suppressing the dependence of the payoff on the co-player's strategy) and it is in accordance with the matrix below.

	P	R
P	a, a	b, c
R	c, b	d, d

We assume that (1) $a > c, d > b$, (2) $a > d$ and (3) $a + b < c + d$. Condition (1) says there are two strict equilibria represented by the strategy pairs (P, P) and (R, R) ; (2) says the first equilibrium payoff dominates the other, while (3) says that the latter is risk dominant. While we follow the literature in focussing on tension between Pareto-efficient coordination and risk-dominant coordination, we only need the weaker condition of $b < c$ instead of the risk-dominance condition (3) above. This weaker condition, which is implied by (3) (but not the other way round), signifies that in the event of mis-coordination, the strategy P user does worse than the strategy R user.⁹ Thus, the results of this chapter are valid for a wider class of games where the restriction is only on the ordinality of the payoffs and where coordinating on strategy R need not be the risk-dominant equilibrium.

After an individual receives the appropriate payoff, he observes the experience (i.e. chosen strategy and payoff received) of a set of other individuals, called his neighbours. The set of neighbours of an individual i , denoted by $N(i)$, is determined by the structure of links in an exogenous network. Individual i is directly connected to all individuals in $N(i)$. It is assumed that $i \notin N(i)$. However, since an individual always knows of his own experience, his reference set is $N(i) \cup \{i\}$. The links are undirected, so that if individual i observes individual j , then individual j also observes individual i . Thus, $j \in N(i)$ implies $i \in N(j)$. Let $N^2(i)$ represent the individuals who are neighbours or neighbours of the neighbours of individual i , i.e. $N^2(i) = \bigcup_{j \in N(i)} N(j) \cup N(i) \setminus \{i\}$. Further, $N^k(i)$ is defined recursively: $N^k(i) = \bigcup_{j \in N^{k-1}(i)} N(j) \cup N^{k-1}(i) \setminus \{i\}$ for $k \geq 2$; it represents the set of other individuals who are at most k links away from individual i . Since we assume the network to be connected, for any individual i , there exists an integer K (for instance, $K = M - 1$) such that $N^K(i) = I \setminus \{i\}$. So, starting from an individual i , every other individual is at a distance of at most K links. The network is said to be complete if, for all individuals $i \in I$, $N(i) = I \setminus \{i\}$; we refer to this as the situation of full observability. Otherwise, we

⁷Under these conditions, with full observability, the state where everyone plays the strategy of Pareto-efficient equilibrium is stochastically stable. See footnote 10 for more details.

⁸For ease of presentation, for most of the chapter, we focus on 2×2 symmetric coordination game. The results generalise to larger games, the arguments for which are put forth in Section 5.4.

⁹Assuming $b > c$ makes matters trivial as there is no force that acts against selection of strategy P .

are in the realm of incomplete observability. The network, N , is defined by the direct links of the individuals: $N = \{N(i)\}_{i \in I}$. The set of all possible connected networks is denoted by \mathcal{N} .

After realising his own payoff and observing the experience of his neighbours, individual i chooses, in the succeeding period, the strategy of the most successful individual he observes. Therefore, he chooses $s_i^* = \operatorname{argmax} u_j(s_j)$, where $j \in N(i) \cup \{i\}$.¹⁰ When the network is complete, $s_i^* = s_j^*$ for any two individuals i and j , i.e. everyone imitates the same strategy. On the other hand, when the network is not complete, then the strategy s_i^* may not coincide with the strategy of the most successful individual in the population. It is also possible that s_i^* and s_j^* are different. Hence different individuals may revise to different strategies under limited observability.

The elements above describe the evolution of behaviour in the society. Individuals receive payoffs depending on their strategy and the choice of strategy of the randomly matched partner. For the following period, they revise their strategy simultaneously and choose the strategy of the most successful individual they know of. This situation continues recurrently under the specified dynamics of “local” imitation.

For a given network and specification of the random matching process, the state of the process at any given point in time is described by the strategy choice of each individual and the state space can be represented by $\Omega = \{P, R\}^M$. It is then possible to define, for a given network and specification of the random matching process, a Markov process Q over the state space that indicate the probabilities of transition between states. Let \vec{P} (\vec{R}) denote the state where all individuals play strategy P (R). In addition, we define a statistic that can be derived from the state of the process: η , the number of individuals using strategy P . For example, when the state is \vec{P} , then $\eta = M$ – everyone uses the strategy of the Pareto-efficient equilibrium. The initial state of the process, and thus the initial description of the derived statistic η , is given arbitrarily.

An absorbing state of the process is defined as a state in which the system remains locked in and there is no possibility of transiting to a different state. Therefore, if the process hits an absorbing state, then the number of individuals who use strategy P (R) does not change. A first observation is that for any network $N \in \mathcal{N}$, the only absorbing states of the process of local imitation are the states \vec{P} and \vec{R} , where all individuals choose the same action. That these two states are absorbing is obvious: if everyone uses the same strategy, then there is no other strategy to imitate and hence the process is locked in thereafter. Further, any other state is non-absorbing; from any other state, the process eventually ends up in either \vec{P} or \vec{R} . This is the content of the lemma below, the proof of which is in the appendix.

¹⁰Note that because the payoffs are distinct, for individual i , s_i^* is unique.

Lemma 5.1. *For any connected network $N \in \mathcal{N}$, the absorbing states of the unperturbed process of local imitation are \vec{P} and \vec{R} , and there are no additional absorbing sets.*

This lemma is immediate for a complete network i.e. under full observability. Everyone imitates the action of the most successful individual, and since payoffs are distinct, all individuals revise to the same strategy, thereby leading the process to one of the absorbing states, \vec{P} or \vec{R} . However, when the network is not complete (i.e. under incomplete observability), it is possible that the random matching unfolds in a way such that $s_i^* = P$ for individual i whereas $s_j^* = R$ for another individual j . Individuals i and j would then imitate different strategies and the transition to the absorbing states would not be immediate; we show this with the help of the following example.

Example. In Figure 5.1 below, the solid lines in the network of ten individuals indicate the links. In the initial state, the individuals numbered one to five use strategy P , while the individuals numbered six to ten choose strategy, R . The dotted lines indicate a particular realisation of the random matching in a period. The triplet alongside each individual indicates, from left to right, his current strategy, payoff received, and strategy he will imitate next period. For example, individual 5 plays strategy P and receives a payoff of b ; in the next period he continues to play strategy P . With the indicated matching, it

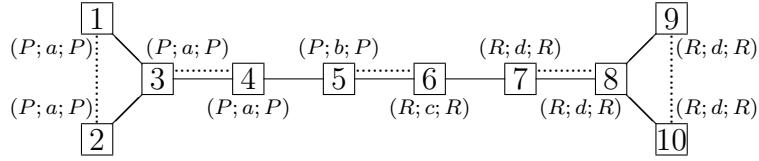


Figure 5.1: An example.

is easily verified that all individuals hold on to their current strategy. The first four (last four) individuals only observe individuals with strategy P (R), and hence do not change their strategy. Individual 5 receives a payoff of b with strategy P and in addition, he observes individual 4 and individual 6 receiving a payoff of a and c with strategies P and R respectively; for him strategy P is the more successful strategy (even though he himself does not fare well with that strategy) and he imitates it in the next period. For similar reasons, individual 6 chooses strategy R in the next period. ■

Thus, in this example, unlike in the full observability model, after a period of strategy revision, the population may consist of both strategy P and strategy R players. However, the transition to \vec{P} or \vec{R} happens eventually. Crucial to the argument is that if the state is other than \vec{P} or \vec{R} , then because of connectedness of the network there exists at least one individual who observes at least one individual who uses strategy P and at least one

individual who uses strategy R . Using this, (we show in the proof in the appendix) it is possible to construct a path in which either the number of individuals who use strategy P , or the number of individuals who use strategy R , increase until all individuals in the population use the same strategy.

The process of pure imitation is now perturbed with the possibility of making mistakes or experimenting; with an independent and time invariant probability ε , an individual plays a strategy randomly. Let Q_ε denote this perturbed Markov process. This process does not have any absorbing states due to the randomness in strategy choice of each individual, and it is now possible to transit from the previously described absorbing states of the unperturbed process Q . Standard results on Markov process indicate that the perturbed process Q_ε has a unique invariant distribution $\mu(\varepsilon)$, i.e. $\mu(\varepsilon)Q_\varepsilon = \mu(\varepsilon)$. To determine the long-run outcome, we use the notion of stochastic stability (see Young (1993, 1998) and Kandori et al. (1993)), which selects the states that receive a positive weight in the limiting invariant distribution, $\mu^* = \lim_{\varepsilon \downarrow 0} \mu(\varepsilon)$. The absorbing state of the unperturbed process that is most resistant in the face of these experimentations is stochastically stable and is expected to be observed more frequently in the long-run, as individuals' probability of making random choices vanishes to zero.

5.3 Results

By Lemma 5.1, the process reaches one of the absorbing states of the unperturbed process in finite time. Let $\rho(\vec{P} \rightarrow \vec{R})$ and $\rho(\vec{R} \rightarrow \vec{P})$ denote the resistance or minimum number of experimentations required to move from the state \vec{P} to the state \vec{R} , and vice-versa, respectively. The state \vec{P} is stochastically stable if $\rho(\vec{P} \rightarrow \vec{R}) \geq \rho(\vec{R} \rightarrow \vec{P})$; further, it is uniquely so if the inequality is strict. Similarly, the state \vec{R} is stochastically stable if $\rho(\vec{P} \rightarrow \vec{R}) \leq \rho(\vec{R} \rightarrow \vec{P})$ and uniquely so if the inequality is strict.

First, consider the benchmark model of full observability where the network is complete. We argue (as in Robson and Vega-Redondo (1996)) that \vec{P} is the unique stochastically stable state. The resistance of the transition from \vec{R} to \vec{P} is equal to 2, i.e. $\rho(\vec{R} \rightarrow \vec{P}) = 2$. To see this, suppose that the initial state is \vec{R} . Experimentation to strategy P by one individual only results in a reversion to the initial state; however, if two individuals experiment with strategy P , and they are matched together (a positive probability event), they realise the highest possible payoff of a . This is observed by everyone and imitation leads to the state \vec{P} . The transition from state \vec{R} to \vec{P} , on the other hand, requires strictly more than two experimentations. To see this, note that when any two individuals experiment with strategy R when the state is \vec{P} , then since there are at least six individuals ($M \geq 6$), there always will be at least two individuals that collect a payoff of a by using

strategy P . So, strategy P is imitated by all individuals and the state reverts to \vec{P} .¹¹ This is summarised in the following remark.

Remark 1. *For a complete network, which means full observability, the state where everyone plays the Pareto-efficient equilibrium strategy is the unique stochastically stable state.*

Now, suppose the connected network is not complete, implying that there is at least one individual who does not observe everyone else. Is the state \vec{P} still stochastically stable? The following proposition (the proof is in the appendix) reveals that it indeed is.

Proposition 5.2. *For any connected network $N \in \mathcal{N}$, the state \vec{P} where everyone uses the Pareto-efficient equilibrium strategy P is a stochastically stable state, while the state \vec{R} where everyone uses the risk-dominant equilibrium strategy may be stochastically stable as well.*

Here, because of limited observability, the arguments made earlier to establish Remark 1 do not apply directly. To illustrate, suppose that starting from the initial state \vec{R} , two individuals experiment with strategy P . With positive probability, they are matched to each other and they receive the highest payoff possible; but due to limited observability, this may not be observed by all and so, all individuals may not imitate the use of strategy P in the next period. However, the transition from \vec{R} to \vec{P} occurs with positive probability following experimentation by two individuals. Connectedness of the network makes this contagion possible and so, $\rho(\vec{R} \rightarrow \vec{P}) = 2$. It can also be shown that the transition in the other direction, i.e. from \vec{P} to \vec{R} needs at least two such experimentations; so, $\rho(\vec{P} \rightarrow \vec{R}) \geq 2$. Since $\rho(\vec{P} \rightarrow \vec{R}) \geq \rho(\vec{R} \rightarrow \vec{P})$, the result follows.

The proposition only indicates that the state where everyone plays the Pareto-efficient strategy is stochastically stable; it does not imply that it is the unique stochastically stable state. In fact, the state \vec{R} may be stochastically stable as well. This possibility arises, and is demonstrated in the example below, when the transition from \vec{P} to \vec{R} is effected with two experimentations.

Example. Figure 5.2 presents an example where the transition from \vec{P} to \vec{R} is affected by two initial experimentations, leading to the conclusion that the state where all individuals use strategy R is also stochastically stable.. The solid lines indicate the links in the network of ten individuals. The initial state in Period 0 is \vec{P} , as indicated by the strategy choice of

¹¹If $M = 2$, the transition from \vec{P} to \vec{R} is possible with one experimentation; further, the transition in the reverse direction take two experimentations. Thus, for $M = 2$, the state \vec{R} is uniquely stochastically stable. If $M = 4$, then a minimum of two experimentations are required to induce the transition from \vec{P} to \vec{R} and from \vec{R} to \vec{P} ; so, both \vec{P} and \vec{R} are stochastically stable states. We assume $M \geq 6$ as it is a more reasonable population size; also, it results in the \vec{P} being the unique stochastically stable state under full observability and makes for meaningful comparisons with the outcomes of incomplete observability.

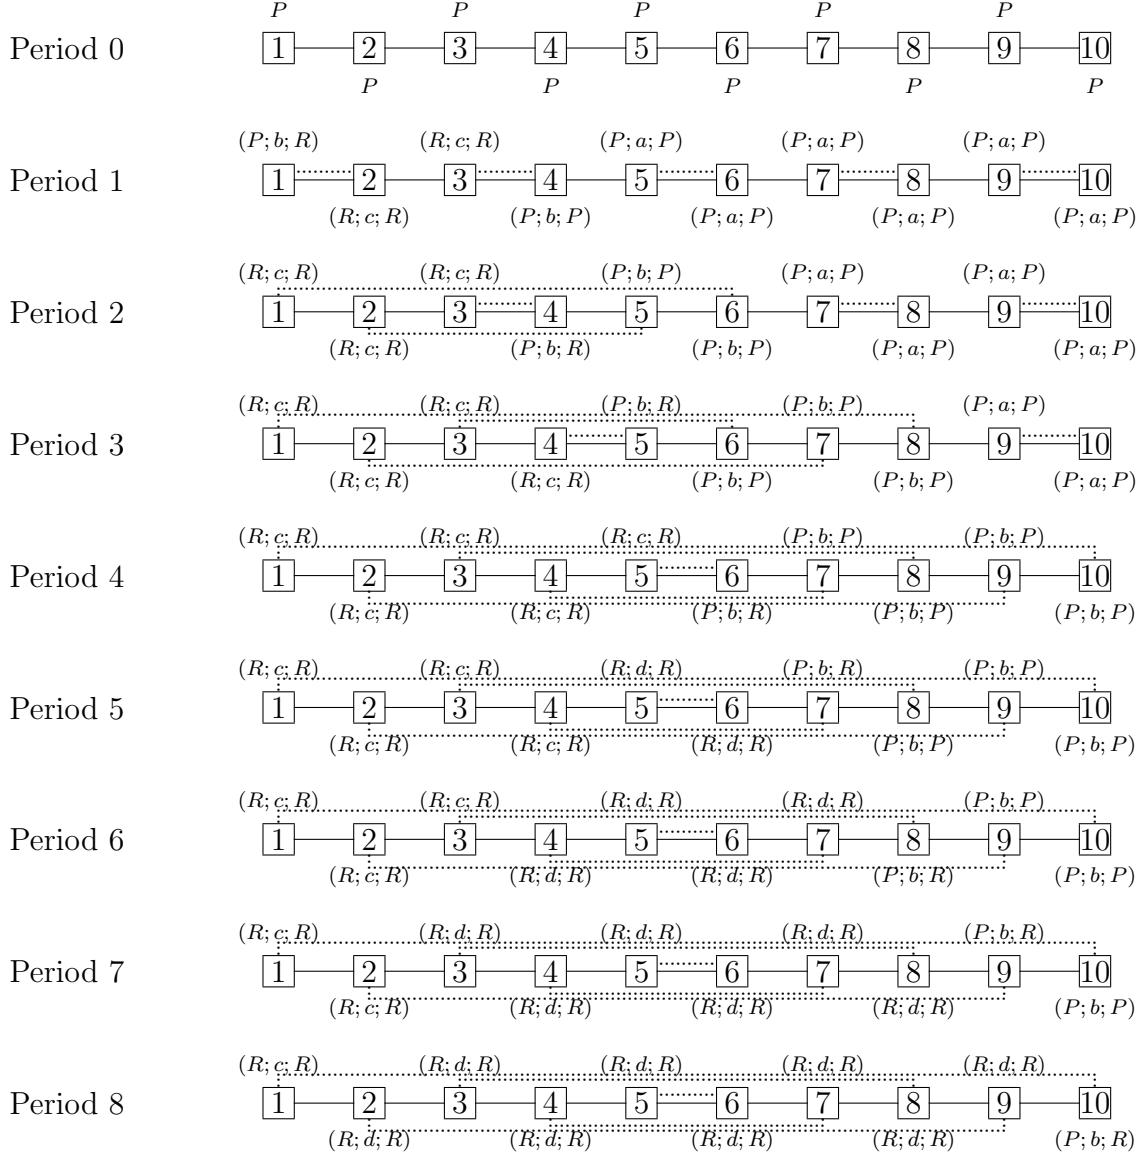


Figure 5.2: Example to illustrate the stochastic stability of state \vec{R} .

each individual in the diagram. Now, suppose that in Period 1, individual 2 and individual 3 experiment with strategy R . The realisation of the random match is depicted with dotted lines: for example, individual 1 is matched to individual 2. The triplet alongside each individual indicates, from left to right, the current strategy, the payoff received, and the strategy to be chosen by imitation in the next period. As a result of the realised matching of individuals, for the first three individuals ($i = 1, 2, 3$), R is the more successful strategy in their respective neighborhoods, while each of the remaining seven individuals ($i = 4, \dots, 10$) observes at least one individual who receives the highest possible payoff of a by application of strategy P . Therefore, at the beginning of Period 2, the first three individuals imitate

strategy R ; all others use strategy P . The next diagram in Figure 5.2 shows a possible realisation of the matching process in Period 2. It can now be verified that, after revision, the first four individuals adopt strategy R . The diagrams for the succeeding periods show the spread of strategy R to the entire population by the local imitation dynamic.

This shows that the state where everyone plays the risk-dominant strategy is also stochastically stable. Hence, limited observability does make an actual difference in the long-run behaviour.¹² ■

Having established that both absorbing states of the unperturbed process of local imitation may be stochastically stable, we try to examine conditions under which the Pareto-efficient outcome will uniquely prevail in the long-run. The stability of the risk-dominant outcome is clearly driven by the fact that an individual is only partially informed. In the example above, individual 1 only observes individual 2 and this feature may be exploited along the least resistant path of transition from \vec{P} to \vec{R} ; had individual 1 been able to observe sufficiently many other individuals, he would have been more resistant to this invasion by strategy R . This observation, in conjunction with Proposition 5.2, suggests the next proposition. The proof is presented in the appendix.

Proposition 5.3. *The state where individuals play the strategy of the Pareto-efficient equilibrium is uniquely stochastically stable in all connected networks where each individual observes at least four other individuals.*

The proposition reveals the need for a minimum amount of connectivity and observability in the network to foist the Pareto-efficient state as the unique stochastically stable state. Further, the condition in the proposition can be interpreted as an absolute condition as it is independent of factors such as population size or the payoffs. Even though the proposition only captures a sufficient condition, and as such does not provide a full characterization, we now show with an example, that the condition is not very loose.

Example. In Figure 5.3, we present a network which marginally deviates from the sufficient condition in Proposition 5.3, and we demonstrate that the transition from \vec{P} to \vec{R} is affected after two initial experimentations, leading to the conclusion that the state \vec{R} is stochastically stable as well. The solid lines in the network indicate the links; only four individuals (marked by the numerals 1, 2, 3, 12) have three links while all other individuals

¹²This contrasts the results derived in Alós-Ferrer and Weidenholzer (2008), where individuals interact with their neighbours but observe and receive information from a wider set of individuals. Individuals choose strategies by imitating the most successful individual they observe. They show that for an appropriately large population, the state where everyone plays P is uniquely stochastically stable. The example here can be extended to show that for a population of any size, the state where everyone plays R may also be stochastically stable.

have four neighbours.¹³

The initial state (Period 0) is \vec{P} . In Period 1, suppose individuals 1 and 3 experiment with strategy R . The realisation of matches is indicated by dotted lines. The triplet next to each individual indicates, from left to right, the current strategy of the individual, the payoff received by the individual as a result of the indicated matching, and the strategy that the individual will revise to by local imitation. The diagrams show the transition to \vec{R} .

Thus, we see that in face of this small deviation, the transition from \vec{P} to \vec{R} is possible with two experimentations and so, the state \vec{R} is also stochastically stable in this network, enabling the conclusion that the sufficient condition in Proposition 5.3 is not very loose. ■

Thus far, no structure is imposed a priori on the shape of the network that governs the flow of information. Now, we turn attention to regular networks. A network is said to be regular when each individual in the population has the same number of links. In terms of our notation, for all individuals $i \in I$, $|N(i)| = k$ for some positive integer $k \leq M - 1$; each individual in the network has k links, where k is the degree of the network.

Proposition 5.4. *For all regular connected networks $N \in \mathcal{N}$, the state where all individuals play the strategy of the Pareto-efficient equilibrium is uniquely stochastically stable.*

This proposition relaxes the minimum observability condition in the previous proposition at the cost of imposing some structure on the network architecture. The proof of this proposition is presented in the appendix, but herein we outline the basic arguments that support the result. It is sufficient to restrict attention regular networks of degree two and three for the reasons that, firstly, for a population size greater than two, there exists no regular connected network with degree one. Secondly, for regular networks of degree higher than three, the statement in the proposition may simply be derived from Proposition 5.3. In regular networks of degree two or three, starting from \vec{P} , if two individuals choose strategy R by the way of experimentation, then the number of individuals who play strategy R in the next period (after revision via local imitation) never exceeds two and, hence, cannot spread to the entire population. For example, if the two experimenting individuals i and j are matched to two other strategy P individuals (represented by k and l), then the only individuals who might consider adopting strategy R by local imitation

¹³It is possible to replicate this network to add more individuals such that only four individuals (the first three individuals and the last individual) have three links while others have four links; thus, the fraction of people who have three links goes to zero as the population size increases, and in this sense, this represents a “small” deviation from the condition in Proposition 5.3.

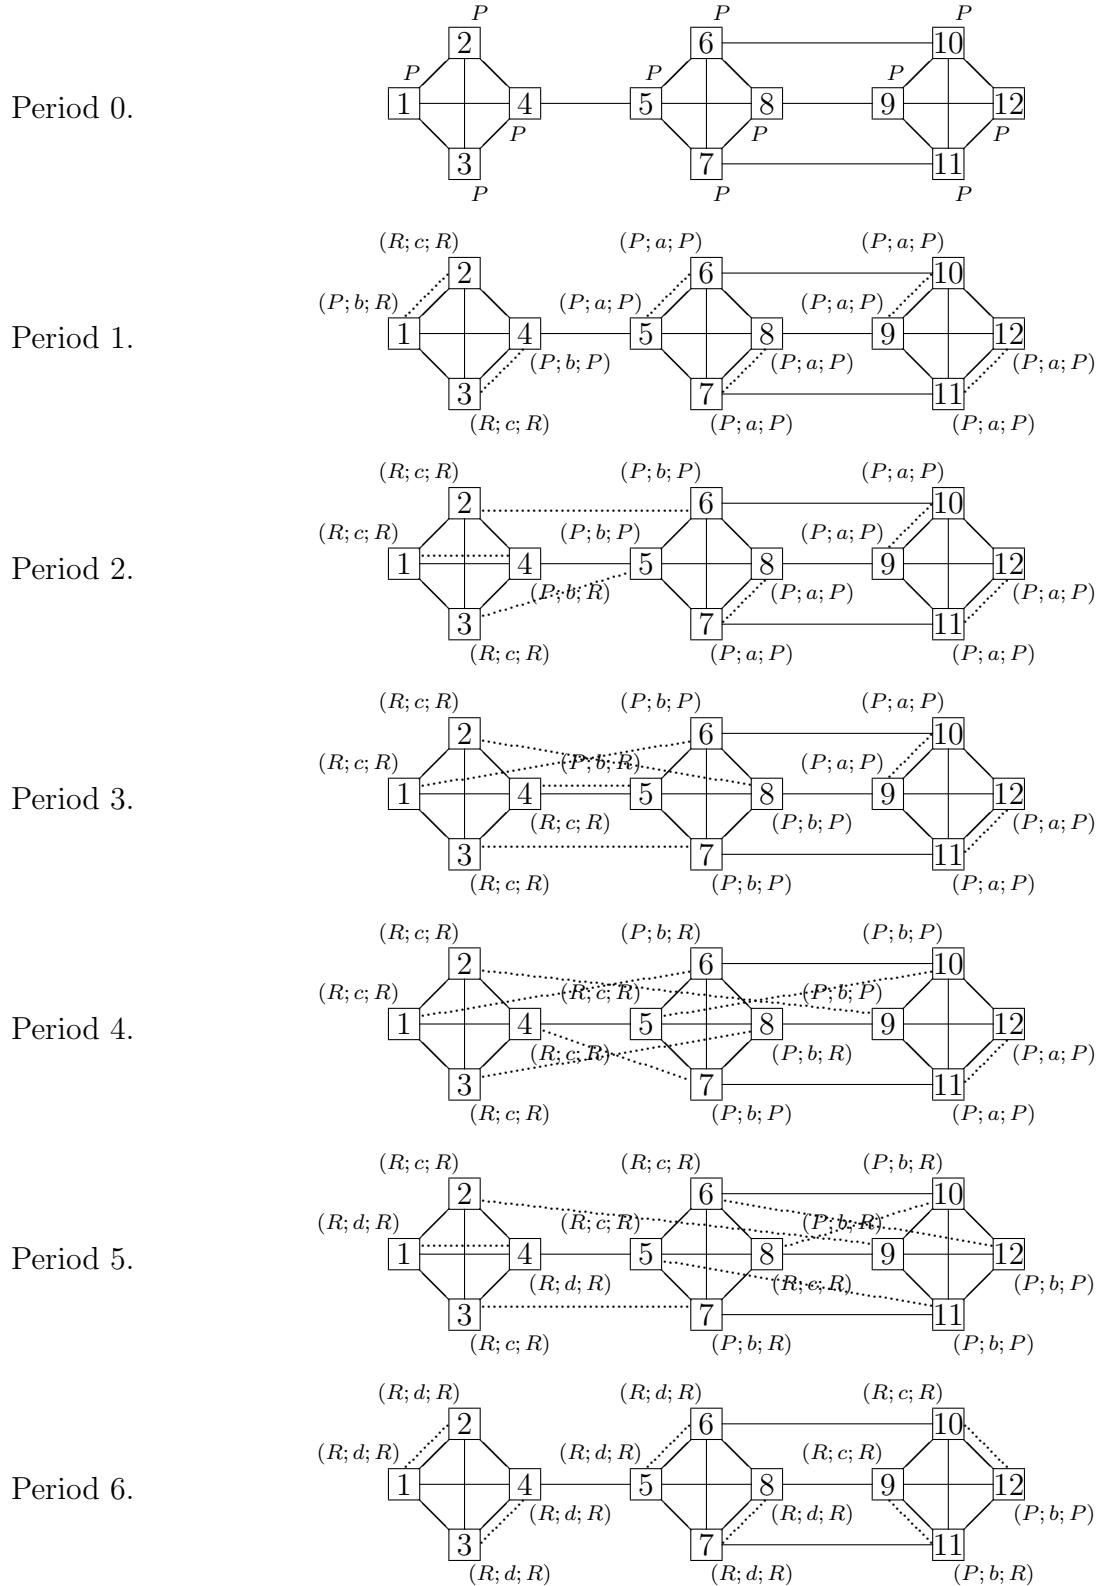


Figure 5.3: Example to show the sufficient condition in Proposition 5.3 is reasonably tight.

are i, j, k and l (as other strategy P individuals realise the highest possible payoff); but, in any regular network, it is not possible for more than two individuals to locally imitate strategy R . It can then be concluded that the transition from \vec{P} to \vec{R} requires strictly more than two experimentations, thereby establishing \vec{P} to be the unique stochastically stable state.

Again, we comment on the slackness of the sufficient condition above. The example associated with Figure 5.2 (which was used to demonstrate the stochastic stability of the risk-dominant equilibrium) represents the “smallest” possible deviation from a regular network of degree two, for the reason that if one keeps on increasing the number of individuals in the network, then only two individuals continue to share only one link¹⁴; thus, the fraction of individuals who do not have exactly two links goes to zero. Yet, in that network, the state where all individuals play strategy R is stochastically stable. This may permit the conclusion that the sufficient condition is reasonably tight.

5.4 Discussion

In this section, we discuss three related aspects of the situation that has been the subject of our study in this chapter. The first issue relates to the observation structure, and we analyse the bearing that observing the performance of the randomly chosen co-player has on the coordination outcome. Thereafter, we discuss the effect of introducing inertia in the imitation dynamic, i.e. what if individuals do not instantaneously imitate the most successful neighbour but do so only with a positive probability? Finally, we close this section by reasoning that the results obtained in this chapter may be extended to larger symmetric coordination games.

Firstly, throughout this chapter, we have assumed that an individual may only imitate his neighbours. This precludes the possibility of him imitating the randomly chosen co-player (if the latter is not a neighbour), when in fact, it may be argued that since the situation is described by a simple coordination game, one may be able to observe (or figure out) and hence, possibly imitate the strategy of the latter. What difference would the consideration of such a possibility make? The answer is simply that all the stated results remain unaffected. Obviously, the states \vec{P} and \vec{R} continue to be the only two absorbing states. As for stochastic stability, a difference might arise only if the minimum resistance of transition from one absorbing state to the other decreases.¹⁵ Firstly, it is

¹⁴Further, it is not possible to have a connected network where all but one individual has degree two, which is why the network in Figure 5.2 represents the smallest possible deviation from a regular connected network of degree two.

¹⁵This consideration might matter only if an individual is matched to a non-neighbour, as only then will the reference set of an individual change. In addition, inclusion of the property that one may observe and

not possible to affect the transition from \vec{R} to \vec{P} with one experimentation. Secondly, $\rho(\vec{P} \rightarrow \vec{R})$ can never equal 1.¹⁶ Hence, Proposition 5.2 holds and the state \vec{P} is always stochastically stable. Similarly, it can be verified that Proposition 5.3 and Proposition 5.4 are unaffected. Furthermore, the examples presented to demonstrate the “tightness” of the sufficient condition continue to be valid.

It is interesting to see that even though Proposition 5.3 involves an minimum amount of observability (of four links), the sufficient condition therein is invariant to the fact that individuals may have an observation in addition to that of their neighbours when they are matched with non-neighbours. In particular, the example associated with Figure 5.3 may be used to realise that for unique stochastic stability of the Pareto-efficient state, the sufficient condition is that each individual should be endowed with at least four links, and that it is not sufficient for individuals to have at least three links even when they obtain additional information from their randomly matched partner. In sum, we can draw the same conclusions even when we open the possibility of an individual observing and imitating his randomly chosen partner, irrespective of them being (or not being) neighbours.

Secondly, we consider a slightly different dynamic by adding the element of inertia. Now, in the unperturbed imitation protocol, individuals imitate the most successful strategy in the neighbourhood with positive probability, and not, as before, with probability equal to one. In case they do not imitate the most successful strategy, they hold on to their most recent strategy. While it is still true (for the same reason stated earlier) that under complete observability, coordinating on the Pareto-efficient equilibrium is the uniquely stochastically stable, it is interesting to note that under incomplete observability, coordination on the risk-dominant equilibrium may be the uniquely stochastically stable state, and we illustrate this by means of the following example.

Consider a star shaped network, where individual 1 is connected to all other individuals, and none of the other individuals share a link amongst themselves. Suppose the initial state is \vec{P} , and individual 1 experiments with strategy R and is matched to individual 2. Individual 2 would observe strategy R to be more successful and imitate it (with positive probability), while it is possible that individual 1 sticks on with strategy R due to inertia. Thus, these two individuals choose strategy R in the next period. Now, individual 1 may

imitate the randomly chosen partner under the imitate-the-best dynamic, cannot increase the minimum resistance of transition.

¹⁶The reason is that if one individual experiments with strategy R when the initial state is \vec{P} , then the only individuals who might choose strategy R by imitation are the experimenting individual and his randomly chosen partner. Both of them would adopt this strategy only if they were not connected to any other strategy P individual, or in other words, if these two individuals were connected only to each other. However, under the assumption of a connected population of at least six individuals, no more than one individual would continue to play strategy R , and hence the transition to \vec{R} is not possible with one experimentation.

be matched to another strategy P individual, who may similarly imitate strategy R (with individual 1 holding on to strategy R due to inertia). In this manner, this strategy may spread throughout the population. Hence, resistance of the transition \vec{P} to \vec{R} equals one while the reverse transition is not possible without two experimentations. So, \vec{R} is the unique stochastically stable state.

The unique stochastic stability of \vec{R} is driven by the fact that there is an individual who only observes one other individual. (Similarly, with inertia in local imitation, \vec{R} is the unique stochastically stable state in the network illustrated in Figure 5.2. This is made possible by individual 1 observing only one other individual.) The spread of strategy R throughout the entire population in these examples, following a single experimentation, hinges on the event that there is an individual (such as individual 2 in the star-shaped network example) who is matched to, and observes, only the experimenting individual. This individual may imitate strategy R , while the experimenting individual plays strategy R in the next period due to inertia. Thus the number of individuals using strategy R grows to two. Clearly, the use of strategy R would not grow beyond the experimenting individual if each individual observed at least two other individuals. Under this restriction of each individual possessing at least two links, both absorbing states may be stochastically stable under incomplete observability (exactly as in Proposition 5.2). Furthermore, if each individual observes four other individuals, then we again re-obtain the result that the Pareto-efficient equilibrium is uniquely stochastically stable, as in Proposition 5.3. The example used earlier in Figure 5.3 is also relevant here and continues to suggest that this sufficient condition is not very slack.

Inertia in local imitation does, however, have a more telling impact on coordination in regular networks. While earlier, Pareto-efficient coordination was the unique stochastically stable state, the risk-dominant equilibrium may now be stochastically stable in regular networks (of degree two and three) when individuals exhibit inertia in local imitation. Even though the example to follow is applicable for any population size, for now, we consider a regular network of degree two (i.e. a circle) with 8 individuals, labeled by natural numbers, with $N(i) = \{i+1, i-1\}, i \in \{1, 2, \dots, 8\}$. (For $i=1$, $N(i) = \{2, 8\}$, and for $i=8$, $N(i) = \{1, 7\}$.) Suppose the initial state is \vec{P} , and individuals 1 and 4 experiment with strategy R , and they are matched to individuals 2 and 3 respectively. Then, the latter two individuals may imitate strategy R , while the former two may hold on to it due to inertia; so, with positive probability, the first four individuals now play strategy R . From here, it is easy to show the spread of this strategy in the entire population. Since $\rho(\vec{R} \rightarrow \vec{P}) = \rho(\vec{P} \rightarrow \vec{R}) = 2$, the risk-dominant equilibrium is stochastically stable along with the Pareto-efficient equilibrium. Examples to the same end may also be constructed

for regular networks of degree three.¹⁷ Therefore, in context of regular networks, inertia only serves to underline the importance of observability, and that now, any weakening of the minimum observability condition may leave the door open for the rooting of inefficient coordination in the population.

Finally, we argue that the results presented generalise to symmetric $m \times m$ coordination games where the coordination outcomes are Pareto-ranked. Suppose that each individual chooses from m strategies (labeled $1, 2, \dots, m$), and for two randomly drawn individuals, coordinating on the same strategy is an equilibrium (and hence, preferable to mis-coordination). We assume that the resulting m equilibria can be Pareto-ranked and without loss of generality, coordinating on strategy u is Pareto-superior to coordinating on strategy v , for all feasible $u < v$. Thus, focussing on the coordinated action (or diagonal) payoffs, coordinating on strategy 1 (strategy m) yields the highest (lowest) payoffs. Let \vec{u} denote the state where all individuals play strategy u (where $u \in \{1, 2, \dots, m\}$). As with 2×2 coordination games, to make matters non-trivial, we assume that on mis-coordination, the individual using the strategy labeled with the lower number does worse than the individual using the strategy with the higher number label (recall footnote 8).

Then, one may realise that: (i) the absorbing states correspond to monomorphic states where all individuals choose the same strategy, (ii) under complete observability, Pareto efficient coordination is achieved,¹⁸ (iii) under incomplete observability, Proposition 5.2 holds, i.e. $\vec{1}$ is stochastically stable, but not uniquely so, (the arguments for the 2×2 coordination game also applies here, as does the example that shows that Pareto-efficient coordination may not be uniquely stochastically stable) (iv) for the same reasons provided earlier, if at least four other individuals are observed, then the most preferred coordination outcome is uniquely stochastically stable, and, (v) for the same reasons provided earlier, if the network is regular, then the most preferred coordination outcome is uniquely stochastically stable. Thus, all the results may be readily generalised to symmetric coordination games with a larger strategy space.

¹⁷Consider a population of size eight, with $N(1) = \{2, 3, 4\}$, $N(2) = \{1, 3, 4\}$, $N(3) = \{1, 2, 5\}$, $N(4) = \{1, 2, 6\}$, $N(5) = \{3, 7, 8\}$, $N(6) = \{4, 7, 8\}$, $N(7) = \{5, 6, 8\}$ and $N(8) = \{5, 6, 7\}$. From the initial state \vec{P} , suppose individuals 3 and 4 experiment with strategy R , and are matched to individuals 1 and 2. With positive probability, the latter two individuals may imitate strategy R , while the former two may hold on to it due to inertia. Thus, four individuals use strategy R , and from here, it is easy to show that with positive probability, strategy R envelops the entire population.

¹⁸The argument for this remains exactly the same – with complete observation, the path of least resistance to $\vec{1}$ from absorbing states equals two, but a transition from $\vec{1}$ takes more than two experimentations. A larger strategy space opens up the possibility that the experimentations may be with different strategies, but even under this consideration, two experimentations do not result in a transition from $\vec{1}$. This is because (owing to the structure of payoffs in a symmetric coordination game) the coordination payoff with strategy 1 is the highest payoff in the payoff matrix – this comment is also applicable for (iv) and (v) to follow.

5.5 Conclusion

This chapter analyses the situation where individuals belonging to a finite population are engaged recurrently in a strategic environment described by a coordination game where the Pareto-efficient equilibrium differs from the risk-dominant equilibrium (or more generally, where the coordination outcomes may be Pareto-ranked). The partner in the coordination game is chosen randomly with full support. The payoffs of the individuals depend on the strategy chosen by him and his randomly drawn partner. In addition to his own experience, an individual is, through his (social) links, aware of the experience of his neighbours. An individual is assumed to imitate the strategy of the most successful experience he is aware of. When all individuals are directly linked (the full observability case), the state where everyone chooses the Pareto-efficient equilibrium action is uniquely stochastically stable. However, this result does not fully carry over when individuals only observe a sub-population. The state where everyone plays the Pareto-efficient equilibrium strategy is always stochastically stable but not uniquely so. For uniqueness of this state in the long run, each individual in the network, independent of the population size or payoffs, has to be endowed with a minimum of four links. The Pareto-efficient state is also uniquely stochastically stable for all regular networks.

The sufficient conditions emphasise the importance of connectedness of individuals in a network-structured society. If all individuals are connected to sufficiently large number of individuals, then one would expect the Pareto-efficient outcome to prevail. However, even the mere presence of a small number of loosely connected individuals makes the population vulnerable to being entrenched in a less desirable outcome (i.e. the risk-dominant equilibrium, or a Pareto-dominated coordination outcome in the more general case). This hints at the benefit that actions such as community building efforts and platforms for interaction amongst individuals confer upon the entire population. Encouraging each individual to develop and share links and therefore information results in the population being more immune to invasion of less desirable practises. It is also important to note that the channel via which sufficient amount of observability supports the preferred equilibrium is not peer pressure or monitoring; the purely informational role of observability is enough to nudge the population to the better outcome.

5.6 Appendix

Proof of Lemma 5.1.

We show that from any initial state other than \vec{P} or \vec{R} , there is a positive probability of reaching the states \vec{P} or \vec{R} . The proof, in three parts, deals with three mutually exclusive situations, where the initial states are such that (i) $0 < \eta \leq \frac{M}{2}$, (ii) $\frac{M}{2} < \eta < M$ and even and (iii) $\frac{M}{2} < \eta < M$ and odd, and shows the convergence to either \vec{P} or \vec{R} . Before beginning with the proof, we remark on an implication of the network being connected: when the state is such that $0 < \eta < M$, there is at least one strategy P individual who is linked to (and so observes) a strategy R individual, and vice-versa (by undirectedness of the links).

Part I. When the initial state is such that $0 < \eta \leq \frac{M}{2}$, at most half (at least half) of the population plays strategy P (R). In the following period, with positive probability, every strategy P individual is matched to a strategy R individual while the remaining strategy R individuals (if any) are matched amongst each other. All strategy P individuals receive a payoff of b while strategy R individuals receive c or d ; in either case, all strategy R individuals receive higher payoffs than strategy P individuals. By the observation made in the beginning of the proof, there exists at least one strategy P individual linked to a strategy R individual. For that individual, R represents the more successful strategy and he switches. On the other hand none of the strategy R individuals are induced to change their strategy. Hence, the count of individuals using strategy R increases by at least one.

Again, in the next period, with positive probability, a similar matching may be realised whereby every strategy P individual is matched to a strategy R individual and the remaining strategy R individuals are matched amongst themselves. This continues till the population contains only strategy R individuals. This proves the lemma for all initial states where $0 < \eta \leq \frac{M}{2}$.

Part II. Now, the initial states are such that $\frac{M}{2} < \eta < M$ and η is even. With positive probability, all strategy P individuals are matched with another strategy P individual. These individuals get the highest payoff of a and so continue to use strategy P . This implies that all strategy R individuals are matched with another strategy R individual and they receive a payoff of d . Again, there exists at least one strategy R individual who observes a strategy P individual receiving a . If the number of such strategy R individuals is even, then the new state is similar in type to the initial state, i.e. $\frac{M}{2} < \eta < M$ and η is even, but the number of strategy P individuals is higher; then, we iterate the argument above till the state reaches \vec{P} or the number of strategy R individuals who imitate strategy P is odd. The latter possibility makes the state such that η is odd; this case is discussed next.

Part III. Finally, suppose that the initial state is such that $\frac{M}{2} < \eta < M$ and η is odd; so, the number of strategy P players is odd. Since the number of individuals in the population is even, this implies that the number of strategy R individuals is also odd. Let the set of strategy P individuals who are linked to at least one strategy R individual be denoted by B . Let the set of strategy P individuals outside of set B be denoted by C ; an individual in set C is a strategy P individual who is linked only to other strategy P individuals.

First assume that set C is non-null. With positive probability the following matching process is realised: an individual in set B is matched to another individual in set B ; if this leaves an individual in set B unmatched (this would happen when the cardinality of this set is odd, and this in particular, subsumes the case where cardinality of set B is one), he is matched with an individual from set C (note the assumption of non-nullity of set C). Further, the remaining individuals in set C are matched to each other – this might leave an individual in set C unmatched (this would happen when the cardinality of B is odd and set C is non-null), in which case, he is matched to a strategy R individual. This last mentioned strategy P individual gets a payoff of b but is linked to at least another strategy P individual; since all other strategy P individuals receive a payoff of a , he imitates strategy P . Thus, following this realisation of random matching, all individuals in set $B \cup C$ observe strategy P as the more successful strategy and choose it in the succeeding period. Further, the strategy R individuals linked to strategy P individuals imitate strategy P . Hence in the new state, the number of strategy P individuals increase; this argument is iterated till the population is solely comprised of strategy P individuals.

Now, suppose that set C is null i.e. each of the strategy P individuals (who are odd in number) are linked to at least one strategy R individual and hence are contained entirely in set B . First suppose that there is one individual (say, i) in set B , who apart from being linked to a strategy R individual, is also linked to another strategy P individual (say, j).¹⁹ In case of the event that there is more than one such individual, let the individual with the least number of links to strategy R individuals be designated as individual i (in case of a tie, the choice of individual i can be made randomly). Then, with positive probability, individual i is matched to a strategy R individual whereas individuals in $B \setminus \{i\}$ (the cardinality of which is even) are matched amongst themselves. Following this, all individuals in set B continue to employ strategy P as: (i) all individuals in that set except individual i receive the highest payoff and (ii) individual i observes individual j doing so. On the other hand, the strategy R individuals linked to individuals in $B \setminus \{i\}$ (by construction of set B and individual i , there is at least one such strategy R individual) imitate strategy P . Again, the number of strategy P individuals increase, and the arguments used thus far can

¹⁹The situation where set C is null and there is no individual in set B who is linked to another individual in set B is discussed in the next paragraph.

be iterated till we reach the state \vec{P} .

Finally suppose that in addition to set C being null, none of the individuals in set B are connected amongst themselves. Then, since the initial state is described by $\eta > \frac{M}{2}$ (i.e. since strategy P individuals outnumber strategy R individuals and since no strategy P individual is connected to another strategy P individual), there must exist at least one strategy R individual, say k , who is connected to more than one strategy P individual. So, assume that k is linked to at least two strategy P individuals, say u and v . With positive probability, all individuals in set $B \setminus \{u\}$ (the cardinality of which is even) are matched amongst themselves and so, receive the highest payoff and continue to play strategy P . Individual u is then matched with a strategy R individual and receives a payoff of b and observes other strategy R individuals receiving either d or c , each of which is larger than b ; consequently, he revises to strategy R . At the same time, individual k switches from strategy R to P on observing individual v 's success with strategy P . (If any strategy R individual is connected to an individual in $B \setminus \{u\}$, he switches to P by imitation.) Thus, in the new state, in no case does the number of strategy P individuals decrease; but now, in this new state, we are back to the situation where at least one strategy P player (individual k) is linked to another strategy P individual (individual v). So, we have a situation where $\eta > \frac{M}{2}$ and an individual in set B is connected to another individual in set B . Hence, we can use the argument in the preceding paragraph to argue that, with positive probability, the number of strategy P individuals increases, and the arguments used thus far can be iterated till we reach the state \vec{P} .

Hence, we have demonstrated that whenever the state is such that $\eta > \frac{M}{2}$ ($\eta \leq \frac{M}{2}$), with positive probability the number of strategy P (R) individuals increases and, therefore, there exists a path to the state \vec{P} (\vec{R}). This establishes the lemma. ■

Proof of Proposition 5.2.

The proof is in two steps. Part I shows that the $\rho(\vec{R} \rightarrow \vec{P}) = 2$. Part II demonstrates that $\rho(\vec{P} \rightarrow \vec{R}) \geq 2$. This establishes the result as $\rho(\vec{R} \rightarrow \vec{P}) \leq \rho(\vec{P} \rightarrow \vec{R})$.

Part I. Let the initial state be \vec{R} , i.e. everyone uses the risk-dominant equilibrium strategy R . Suppose two individuals i and j – to simplify matters, assume that they are neighbours – experiment with strategy P . With positive probability, they are also matched to each other and they realise the highest payoff of a . This is observed by the neighbours of i and j . Consequently, the individuals who use strategy P in the next period are $N(i) \cup N(j)$ (note that individuals i and j are included in this set). Suppose that this does not make up the entire population. Then there are two possibilities corresponding to the cardinality of the set $N(i) \cup N(j)$ being odd or even. If it is odd, then with positive probability, individuals in the set $N(i) \cup N(j) \setminus \{i\}$ (the cardinality of this set is even) find

their match in the same set, while individual i is matched with a strategy R individual. As individuals in $N(i) \cup N(j) \setminus \{i\}$ themselves realise the highest payoff, they continue to use strategy P while individual i does so because he observes individual j . But now, in addition, the neighbours of individuals in $N(i) \cup N(j)$ also imitate this strategy. Hence, the set of individuals who now use strategy P expands to $N^2(i) \cup N^2(j)$.

On the other hand, if the cardinality of the set $N(i) \cup N(j)$ is even, then with positive probability, each individual in this set is matched to another individual from this set. They receive the highest payoff and their strategy is imitated by their neighbours. Again, the set of individuals who now use strategy P expands to $N^2(i) \cup N^2(j)$. Now, the same reasoning can be used to conclude that the set of individuals who use strategy P expands to $N^K(i) \cup N^K(j) = I$ (where K is the farthest distance between two individuals in the population): if for some $k \geq 2$, the cardinality of $N^k(i) \cup N^k(j)$ is even, then with positive probability, all individuals in that set are matched to another individual from the set; if the cardinality is odd, then all individuals except individual i are matched to another individual from the set. In any case, by local imitation, the strategy spreads to $N^{k+1}(i) \cup N^{k+1}(j)$ until it encompasses the entire population. This shows $\rho(\vec{R} \rightarrow \vec{P}) \leq 2$. It is trivial to see that $\rho(\vec{R} \rightarrow \vec{P}) > 1$; so, $\rho(\vec{R} \rightarrow \vec{P}) = 2$.

Part II. Suppose that from initial state \vec{P} , where everyone uses strategy P , individual i experiments to strategy R . This individual is matched to a strategy P individual, j . Individual i (j) receives a payoff of c (b). By the assumption that the population consists of at least six individuals and by connectedness of the network, it must be true that i or j (if not both) observe a strategy P individual receiving a payoff of a , and so chooses strategy P in the next period. Then, in the next period, at most one individual still chooses strategy R . Thus, it is not possible for strategy R to grow in the population, leading to the conclusion that the transition from \vec{P} to \vec{R} requires at least two experimentations. ■

Proof of Proposition 5.3.

By Proposition 5.2, the state where everyone chooses strategy P is always stochastically stable. Therefore, if we can show that the transition from \vec{P} to \vec{R} takes more than two experimentations when each individual observes at least four other individuals, we can conclude that the state where everyone chooses strategy R is not stochastically stable, and this establishes the proposition.

Let the initial state be \vec{P} , i.e. everyone uses strategy P . Suppose that two individuals, i and j , experiment with strategy R . Irrespective of the realisation of the random matching process, each individual observes at least one individual with strategy P receiving the highest payoff a . To see this, note that there are two types of matches that may be formed after i and j experiment with strategy R :

(a) i and j are matched to each other and so, receive d . All other individuals play strategy P and receive a ; so, none of them switch to strategy R . On the other hand, i and j both observe at least one individual receiving the highest payoff a with strategy P and so, they switch to strategy P . The state reverts to \vec{P} .

(b) i and j are matched to two individuals k and l respectively, who use strategy P . Consequently, i and j receive c ; k and l receive b . However, since each individual is linked to at least four other individuals, each of them observes at least one individual who receives the highest payoff a with strategy P (since $M \geq 6$, such an individual exists). So, they choose strategy P in the next period. The other individuals in the population continue to choose P as they themselves receive a payoff of a with that strategy. Again, the state reverts to \vec{P} .

Thus, the transition from \vec{P} to \vec{R} requires strictly more than two experimentations. ■

Proof of Proposition 5.4.

It will be argued that when the network is regular, and two individuals experiment with strategy R from the initial state \vec{P} , then the number of individuals who choose strategy R in the succeeding periods will never exceed two; hence, the transition from \vec{P} to \vec{R} is not possible with two experimentations. This would prove that \vec{P} is the only stochastically stable state.

Assume that two individuals, i and j , experiment with strategy R . If they are matched to each other, then they receive a payoff of d while everyone else receives a . Hence, the only individuals who might consider playing strategy R in the succeeding period are i and j . However, since they are linked to two or three other individuals (depending on the degree of the regular network), they always observe another individual receiving a payoff of a , and imitate strategy P . It follows that if there is any chance for the population to transit to \vec{R} after the two individuals experiment with strategy R , then they must be matched with strategy P individuals and not with themselves. Resultantly, the focus is on situation where experimenting individuals, i and j , are matched to two strategy P individuals, k and l , respectively. Now, the only individuals who might consider imitating the use of strategy R are $\{i, j, k, l\}$ as all others end up receiving the highest payoff a (recall that we deal with populations containing at least six individuals). It will be shown that in the next period, of these four individuals, no more than two individuals will consider using strategy R and hence the use of this strategy cannot spread to the entire population. Part I shows this for regular networks of degree two while Part II does the same for regular networks of degree three.

Part I. An individual belonging to the set $D = \{i, j, k, l\}$ will choose strategy R by imitation only if he observes strategy R as the most successful strategy in his neighbour-

hood. This in turn happens only if he shares both his links with two other individuals from the set D . Otherwise, he is linked to a strategy P individual who receives the payoff of a , and is led to imitate strategy P . So, let us suppose that two individuals, without loss of generality, i and j , have both their neighbours in set D , thereby implying that both of them choose strategy R in the next period. It will be seen that, given this, it is not possible for either k or l to share both their links with individuals from the set D . Hence, both k and l observe at least one individual outside the set D and they imitate strategy P , thereby making the spread of strategy R impossible.

A connected regular network of degree two can only be a Hamiltonian cycle. Further, in order for both individuals i and j to have two neighbours from set D , it must be that they observe each other. Let the other neighbour of individual i (j) is k (l). If the network has at least six individuals, then it follows that individuals k and l must have a neighbour who receives the highest payoff from the use of strategy P . So, individuals k and l imitate strategy P , and in the next period, no more than two individuals use strategy R . The following paragraphs demonstrate this in more detail.

First, assume that i and j do not observe each other, i.e. $j \notin N(i)$ and $i \notin N(j)$. Then it must be that $N(i) = N(j) = \{k, l\}$ (by the assumption that i and j observe two individuals from the set $\{i, j, k, l\}$). But then this means that $N(k) = N(l) = \{i, j\}$ – a contradiction to the assumption that the regular network of degree two consists of at least six individuals and is connected. (Note that it is not possible to include a fifth individual as i, j, k, l already have exhausted both their links.)

So, let us assume that i and j do observe each other, i.e. $j \in N(i)$ and $i \in N(j)$. Then there are two possible cases:

(i) i and j observe a common individual from the set D , say k . So, $N(i) = \{j, k\}$, $N(j) = \{i, k\}$ and $N(k) = \{i, j\}$ – a contradiction to the assumption that the regular network of degree two consists of at least six individuals, and is connected. (Note that it is not possible to include a fourth individual as i, j, k already have exhausted both their links.)

(ii) i and j observe different individuals from the set D , say k and l respectively. So, $N(i) = \{j, k\}$, $N(j) = \{i, l\}$, $i \in N(k)$ and $j \in N(l)$. As mentioned earlier, in order for individuals k and l to choose strategy R in the next period, it must be that they observe two individuals from the set D . The only such possibility is when $N(k) = \{i, l\}$ and $N(l) = \{j, k\}$ – but again, this contradicts the assumption that the regular network of degree two consists of at least six individuals.

This proves that the use of strategy R does not spread beyond two individuals in the population in a regular network of degree two that contains at least six individuals and so, the transition from \vec{P} to \vec{R} needs more than two experimentations.

Part II. Now, each individual observes three other individuals. For an individual in $D = \{i, j, k, l\}$ to choose strategy R by local imitation after the initial experimentation by i and j , it has to be that he observes the other three individuals in this set. Again, assume that two individuals, say i and j observe three individuals from the above set; in this event, it will be seen that it is not possible for k and l to choose strategy R , thereby limiting the use of this strategy to two individuals (at most). By assumption, $N(i) = \{j, k, l\}$ and $N(j) = \{i, k, l\}$; so, $\{i, j\} \subset N(k)$ and $\{i, j\} \subset N(l)$. If k is to choose strategy R by imitation, then, it must be (by undirectedness of links) that $N(k) = \{i, j, l\}$ and therefore, $N(l) = \{i, j, k\}$. But then this contradicts the assumption that the regular network of degree three has at least six individuals, and is connected. Hence, it is not possible for more than two individuals to choose strategy R by local imitation. This again proves that the transition from \vec{P} to \vec{R} needs more than two experimentations. ■

References

1. Alchian A (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy* 58 (3): 211-221.
2. Alós-Ferrer C, A Ania and KR Schenk-Hoppé (2000). An evolutionary model of Bertrand oligopoly. *Games and Economic Behavior* 33 (1): 1-19.
3. Alós-Ferrer C and S. Weidenholzer (2006). Imitation, Local Interactions and Efficiency. *Economic Letters* 93 (2): 163-168.
4. Alós-Ferrer C and S Weidenholzer (2008). Contagion and Efficiency. *Journal of Economic Theory* 143 (1): 251-274.
5. Alós-Ferrer C and G Kirchsteiger (2010). General equilibrium and the emergence of (non) market clearing trading institutions. *Economic Theory* 44 (3): 339-360.
6. Alós-Ferrer C and S Weidenholzer (2010). Imitation and the Role of Information in Overcoming Coordination Failures. Mimeo, University of Vienna.
7. Apesteguia J and R Selten (2005). Experimentally observed imitation and cooperation in price competition on the circle. *Games and Economic Behaviour* 51 (1): 171-192.
8. Apesteguia J, S Huck and J Oechssler (2007). Imitation: Theory and experimental evidence. *Journal of Economic Theory* 136 (1): 217-235.
9. Apesteguia J, S Huck, J Oechssler and S Weidenholzer (2010). Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust. *Journal of Economic Theory* 145 (5): 1603-1617.
10. Banerjee A and D Fudenberg (2004). Word of Mouth Learning. *Games and Economic Behavior* 46 (1): 1-22.
11. Basu K and J Weibull (1991). Strategy subsets closed under rational behavior. *Economic Letters* 36 (2): 141-146.

12. Basu K and J Weibull (2003). Punctuality - A Cultural Trait as an Equilibrium. In: R. Arnott et al. (Eds.), *Economics for an Imperfect World: Essays in Honor of Joseph Stiglitz*, MIT Press.
13. Bergin J and Bernhardt D (2004). Comparative Learning Dynamics. *International Economic Review* 45 (2): 431-465.
14. Bergin J and B Lipman (1996). Evolution with state-dependent mutations. *Econometrica* 64 (4): 943-956.
15. Blume L (2003). How noise matters. *Games and Economic Behavior* 44 (2): 251-271.
16. Binmore K (1987). Modeling rational players I. *Economics and Philosophy* 3: 179-214.
17. Binmore K (1988). Modeling rational players II. *Economics and Philosophy* 4: 9-55.
18. Björnerstedt J and J Weibull (1996). Nash Equilibrium and Evolution by Imitation. In: K. Arrow, E. Colombatto, M. Perlman, C. Schimdt (Eds.) *The Rational Foundations of Economic Behaviour*. Macmillan Press.
19. Boncinelli L (2008). Global vs. local information in (anti)-coordination problems with imitators. *B.E. Journal of Theoretical Economics (Topics)* 8 (1).
20. Brown GW (1951). Iterative Solutions of Games by Fictitious Play. In: T.C. Koopmans (Ed.) *Activity Analysis of Production and Allocation*. New York: Wiley.
21. Buckner RL and DC Carroll (2007). Self-projection and the brain. *Trends in Cognitive Sciences* 11 (2): 49-57.
22. Camerer CF, TH Ho and JK Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119 (3): 861-898.
23. Chen H-C, Y Chow and L-C Wu (Forthcoming). Imitation, local interaction, and coordination. *International Journal of Game Theory*.
24. Chen Y and MH Riordan (2008). Price-increasing competition. *Rand Journal of Economics* 39 (4): 1042-1058.
25. Coricelli G and R Nagel (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences* 106 (23): 9163-9168.

26. Crawford VP and N Iribarri (2007). Fatal attraction: Salience, naivete, and sophistication in experimental “hide-and-seek” games. *American Economic Review* 97 (5): 1731-1750.
27. Dawkins R (1976). *Selfish Gene*. Oxford University Press.
28. Ellison G and D Fudenberg (1993). Rules of Thumb for Social Learning. *Journal of Political Economy* 101 (4): 612-643.
29. Ellison G and D Fudenberg (1995). Word-of-Mouth Communication and Social Learning. *The Quarterly Journal of Economics* 110 (1): 93-125.
30. Ellison G (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies* 67 (1): 17-45.
31. Eshel I, L Samuelson and A Shaked (1998). Altruists, Egoists, and Hooligans in a Local Interaction Model. *The American Economic Review* 88 (1): 157-179.
32. Goyal S and F Vega-Redondo (2005). Network Formation and Social Learning. *Games and Economic Behaviour* 50 (2): 178-207.
33. Hamilton W (1970). Selfish and spiteful behavior in an evolutionary model. *Nature* 228 (5277): 1218-1220.
34. Hedlund J (2012). Altruism and Local Interaction. *The B.E. Journal of Theoretical Economics (Contributions)* 12 (1).
35. Hehenkamp B and OM Kaarbøe (2008). Imitators and optimisers in a changing environment. *Journal of Economic Dynamics and Control* 32 (5): 1357-1380.
36. Hehenkamp B and Wambach A (2010). Survival at the center – The stability of minimum differentiation. *Journal of Economic Behavior and Organization* 76 (3): 853-858.
37. Herings P (1997). Endogenously determined price rigidities. *Economic Theory* 9 (3): 471-498.
38. Huck S, HT Normann and J Oechssler (1999). Learning in Cournot oligopoly – An experiment. *Economic Journal* 109 (454): C80-C95.
39. Huck S, HT Normann and J Oechssler (2000). Does information about competitors actions increase or decrease competition in experimental oligopoly markets? *International Journal of Industrial Organization* 18 (1): 39-57.

40. Hurkens S (1995) Learning by forgetful players. *Games and Economic Behaviour.* 11 (2): 304-329.
41. Josephson J (2009). Stochastic adaption in finite games played by heterogeneous populations. *Journal of Economic Dynamics and Control* 33 (8): 1543-1554.
42. Kandori M, GJ Mailath and R Rob (1993). Learning, mutation, and long-run equilibria in games. *Econometrica* 61 (1): 29-56.
43. Khan A. (2013). Coordination under Global Random Interaction and Local Imitation. *Meteor Research Memorandum* 13/004
44. Khan A. and R. Peeters (2011). Evolution of Behaviour when Duoplists choose Price and Quantity. *Meteor Research Memorandum* 11/027.
45. Khan A. and R. Peeters (2012). Cognitive hierarchies in adaptive play. *Meteor Research Memorandum* 12/007.
46. Khan A. and R. Peeters (2013). Imitation by Price and Quantity setting firms in a Differentiated Market. *Meteor Research Memorandum* 13/022.
47. Matros A (2003). Clever agents in adaptive learning. *Journal of Economic Theory* 111 (1): 110-124.
48. Mengel F (2009). Conformism and Cooperation in a Local Interaction Model. *Journal of Evolutionary Economics* 19 (3): 397-415
49. Mohlin E (2012). Evolution of theories of mind. *Games and Economic Behavior* 75 (1): 299-312.
50. Nagel R (1995). Unraveling in guessing games: An experimental study. *American Economic Review* 85 (5): 1313-1326.
51. Offerman T, J Potters and J Sonnemans (2002). Imitation and belief learning in an oligopoly experiment. *Review of Economic Studies* 69 (4): 973-997.
52. Robson A and F Vega-Redondo (1996). Efficient Equilibrium Selection in Evolutionary Games with Random Matching. *Journal of Economic Theory* 70 (1): 65-92.
53. Rosenthal RW (1980). A model in which increase in the number of sellers leads to a higher price. *Econometrica* 48 (6): 1575-1579.

54. Ross L, D Greene and P House. (1977). The ‘false consensus effect’: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13 (3): 279-301.
55. Sáez-Martí M and J Weibull (1999). Clever agents in Young’s evolutionary bargaining model. *Journal of Economic Theory* 86 (2): 268-279.
56. Sandholm WH (2010). Population Games and Evolutionary Dynamics. MIT Press.
57. Satterthwaite MA (1979). Consumer information, equilibrium industry price, and the number of seller. *Bell Journal of Economics* 10 (2): 483-502.
58. Schipper B (2009). Imitators and optimisers in Cournot oligopoly. *Journal of Economic Dynamics and Control* 33 (12) : 1981-1990.
59. Stiglitz J (1987). Competition and number of firms in a market: are duopolies more competitive than atomistic markets?. *Journal of Political Economy* 95 (5): 1041-1061.
60. Stahl DO (1993). Evolution of $smart_n$ players. *Games and Economic Behaviour* 5 (4): 604-617.
61. Stahl DO and PW Wilson (1994). Experimental evidence on players’ models of other players. *Journal of Economic Behavior & Organization* 25 (3): 309-327.
62. Stahl DO and PS Wilson (1995). On players’ models of other players: Theory and experimental evidence. *Games and Economic Behaviour* 10 (1): 218-254.
63. Tanaka Y (1999). Long run equilibria in an asymmetric oligopoly. *Economic Theory* 14 (3): 705-715.
64. Tanaka Y (2000). Stochastically stable states in an oligopoly with differentiated goods: equivalence of price and quantity strategies. *Journal of Mathematical Economics* 34 (2): 235-253.
65. Tanaka Y (2001). Evolution to equilibrium in an asymmetric oligopoly with differentiated goods. *International Journal of Industrial Organisation* 19 (9): 1423-1440.
66. Tirole J (1988). The Theory of Industrial Organization. MIT Press.
67. Ullman-Margalit E (1978). Invisible Hand Explanations. *Synthese* 39: 263-291.
68. Van Damme E and J Weibull (2002). Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory* 106 (2): 296-315.

69. Van den Berg, A and Bos, I (2012). Collusion in a price-quantity oligopoly. Meteor Research Memoranda 11/039.
70. Vega-Redondo F (1997). The evolution of Walrasian behavior. *Econometrica* 65 (2): 275-284.
71. Wang J T-Y, M Spezio and CF Camerer (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review* 100 (3): 984-1007.
72. Weidenholzer S (2010). Coordination Games and Local Interactions: A Survey of the Game Theoretic Literature. *Games* 1 (4): 551-585.
73. Young HP (1993a). The evolution of conventions. *Econometrica* 61 (1): 57-84.
74. Young HP (1993b). An evolutionary model of bargaining. *Journal of Economic Theory* 59 (1): 145-168.
75. Young HP (1998). Individual Strategy and Social Conventions. Princeton University Press.

Nederlands Samenvatting

De hoofdstukken in dit proefschrift bestuderen de lange termijn resultaten van strategische situaties wanneer individuen eenvoudige gedragsregels gebruiken. Zoals in hoofdstuk 1 is besproken kan het onderzoeksprogramma dat hier gevolgd wordt geklassificeerd worden onder de brede noemer van evolutionaire speltheorie en beoogt het de meer traditionele speltheorie aan te vullen alwaar rationaliteit en optimaliteit van economische beslissingen de dominante paradigma zijn. Er zijn twee logische vragen die opkomen als men de notie van rationaliteit in keuzegedrag probeert te verzwakken. Ten eerste, in welke mate kan rationaliteit verzwakt worden zonder kwalitatieve verandering van uitkomsten en voorspellingen van het traditionele paradigma?; ten tweede, wat zijn de interessante kenmerken die opduiken wanneer verzwakking van rationaliteit de uitkomsten en voorspellingen veranderd. In dit proefschrift onderzoek ik het effect van het verzwakken van rationaliteit door gebruik te maken van een beste-antwoord dynamiek en een imitatie dynamiek.

In zijn eenvoudigste vorm, in een beste-antwoord dynamiek, speelt de beslisser een beste antwoord op geobserveerde acties. In hoofdstuk 2, gebruikmakend van het concept van beste-antwoord, definieer ik een cognitief hirarchie model. Het meest elementaire niveau van gedrag wordt gegeven door eenvoudig beste antwoord. Het volgende niveau van verfijning (of cognitie) is gedefinieerd door een beste antwoord op een beste antwoord te spelen. Dit kan worden uitgebreid tot hogere niveaus van verfijning in gedrag. Gebruikmakend van zo'n genestelde structuur, is het mogelijk om een model van cognitieve hirarchien te conceptualiseren. Ik constateer dat op populatie niveau, voorbij een zeker niveau van verfijning, een verdere verhoging in niveau van verfijning geen toegevoegde waarde heeft. In hoofdstuk 3 onderzoek ik een (gedifferentieerde) markt met twee bedrijven. Het gedrag van elk bedrijf wordt gedreven door imitatie van het meer succesvolle bedrijf. Het is bekend uit bestaande literatuur dat imitatie vaak leidt tot sub-optimale resultaten terwijl kwaadwillige acties een belangrijke rol spelen in het bepalen van het resultaat van door imitatie gedreven processen. Om een vergelijkbare reden vind ik dat onder sommige omstandigheden imitatie kan leiden tot prijzen die hoger zijn dan wat als optimaal kan worden beschouwd (in de traditionele zin), terwijl onder sommige omstandigheden de prijs aangenomen door elk bedrijf lager kan zijn dan wat traditioneel als optimaal kan worden

beschouwd. Hoofdstuk 4 breidt dit uit naar een markt met een groter aantal bedrijven en hier is de belangrijkste conclusie dat, zelfs met imitatie, het resultaat samenvalt met wat men zou verwachten bij "rationele" bedrijven.

In hoofdstuk 5 onderzoek ik imitatie gedrag in een situatie waarbij individuen gebaat zijn hun acties te coordineren. Terwijl individuen met ieder ander in interactie kunnen zijn, observeren ze slechts de gekozen acties en resultaten van hun sociale groep en nemen ze de acties van het meest succesvolle individu in hun respectievelijke sociale groep aan. Dit doelt om situaties waarbij individuen met elkaar in interactie zijn onder enige mate van onzekerheid en de ervaringen van hun sociale groep gebruiken als richtlijn voor toekomstige acties. Ondanks dat individuen hun acties willen coordineren zijn niet alle vormen van coordinatie gelijk; een bepaalde mate van coordinatie kan geprefereerd worden over andere types of coordinatie. De vraag is of individuen in staat zijn om te coordineren op het meest gewenste resultaat met het oog op de onzekerheid van het niet weten wat de partner in interactie zal doen. Ik constateer dat terwijl deze onzekerheid coordinatie op het meest gewenste resultaat moeilijk maakt, een minimum grootte van de sociale groep coordinatie op het meest gewenste resultaat mogelijk kan maken.

Short Curriculum Vitae

Abhimanyu Khan was born on June 25, 1984. After completing his schooling in Kolkata, he pursued a bachelor's degree in Economics at St. Stephens' College, Delhi University. Thereafter, he obtained a Masters degree in Quantitative Economics from the Indian Statistical Institute, Kolkata and a Masters degree in Economics from the University of Illinois at Urbana-Champaign. He joined Maastricht University in 2010 as a doctoral candidate. He will join as a post-doctoral researcher at the Institute of New Economic Thinking (INET), Cambridge starting September 2014.