

CHAPTER 9

General conclusion and synthesis

Everywhere we look in nature we see cooperation. Yet, the origins and persistence of stable cooperation in the face of potential defection often remains difficult to explain. How did nature solve the problem of cooperation without the rules, norms and authority found in human societies? While we have a good understanding of many cooperative behaviours within species, usually through some form of kin selection, and while mutualisms with (strict) vertical transmission are well understood from an evolutionary perspective too (West *et al.*, 2007b; Bright & Bulgheresi, 2010; Leigh, 2010; Gardner *et al.*, 2011), the evolution of (partially) horizontally transmitted mutualisms is one of the greatest challenges in cooperation evolution. In this thesis, I have expanded our knowledge on the evolution, persistence and functioning of stable cooperation in such systems, by studying belowground cooperation between plants and microbial mutualists. This is important because even though these cooperative interactions represent some of the most ecologically and evolutionarily vital mutualisms (Peoples *et al.*, 1995; Hoeksema *et al.*, 2010; Berendsen *et al.*, 2012; Batterman *et al.*, 2013a), they are less often considered as models for the evolution of cooperation.

Microbial and underground biological markets?

Active selection of cooperative partners, for instance through preferential rewarding of cooperative individuals, can counteract the incentive to defect from cooperation (Sachs *et al.*, 2004; Leigh, 2010). This has the effect of stabilising mutualisms and can result in ‘biological markets’ where individuals compete by offering high-quality mutualistic services (Noë & Hammerstein, 1994, 1995; Leimar & Hammerstein, 2010). It is intuitive to see how this would work in human economies, and relatively easy in animal societies where individuals have advanced cognition that allows them to keep track of exchanges, compare prices and respond in an economic fashion (Fruteau *et al.*, 2009, 2011b; Adam, 2010; Raihani *et al.*, 2012). For biological markets to occur, however, the ultimate requirements are immediate rewards of cooperative behaviour and evaluating the value of a mutualistic benefit in the local environmental context (supply and demand for the service or resource) (Noë, 2006b; Brosnan *et al.*, 2010). Cognition may be helpful in accommodating this, and presumably can produce more advanced (biological) markets including phenomena like future planning and third-party interactions but as I argued in Chapter 2 through evolved responses, market phenomena can in principle be produced even by microbes (Werner *et al.*, 2014c). For instance chemotactic responses of free-living mutualistic microbes might effectively amount to a ‘market choice’ if they result in microbes migrating preferentially to and cooperating with higher quality partners.

The approach of applying principles from human economics to microbes shares a parallel with the origins of evolutionary game theory as a way to understand social interactions (*e.g.* Maynard Smith & Price, 1973). Evolutionary game theory was originally based

on game theory models derived for and applied to humans, such as the famous Nash-equilibrium (Nash, 1950, 1951). While initially it was often (implicitly) framed as of relevance only for actions of conscious and rational actors, this modelling strategy proved useful also in an evolutionary context where responses can arise over evolutionary time through natural selection (Maynard Smith, 1982; Weibull, 1995). The effort to further extend biological market theory from cognitive organisms into the microbial world stands within this wider intellectual tradition of applying analytical tools and concepts from the human domain to biology.

But, has it taught us anything about the actual biology of (microbial) mutualism? Currently, it is too early to tell. The big open question for the field is not so much if microbial biological markets could exist in principle, but rather if partner selection and markets are a common or important explanation for cooperative behaviour involving microbes. Is individual variation in partner quality common among microbes? Can microbes and hosts typically detect it, and select partners based on quality? How do (microbial) partner preferences change in variable environments? Arguably, the only microbial mutualism where we are reasonably certain that the requirements for biological markets are met is the mycorrhizal-plant mutualism (Bever *et al.*, 2009; Kiers *et al.*, 2011). In Chapter 2, I argue that in addition many more microbial cooperative behaviours are at least consistent with its principles and catalogue some potential examples. Now it remains to be seen to what extent detection of partner selection by or of microbes, similar to what happens in arbuscular mycorrhizal fungi, is actually prevalent. Key tools to study this will come from those areas of microbiology that examine (phenotypic) heterogeneity among individual microbial cells (Ackermann, 2013, 2015), as well as from modern tools to track resource flows at fine spatial scales (Whiteside *et al.*, 2009; Hol & Dekker, 2014) and from experimental evolution approaches (Buckling *et al.*, 2009; Hillesland & Stahl, 2010; Harcombe, 2010; Hillesland *et al.*, 2014).

Partner selection and other models of mutualism stability

My primary aim with publishing Chapter 2 of this thesis has been to call attention to this class of market models and to stimulate researchers studying microbial mutualisms to take the partner selection framework into account when analysing cooperative interactions. So far, this seems to have had some success (Franklin *et al.*, 2014; Kemen, 2014; Ponce-Soto *et al.*, 2015; Enyeart *et al.*, 2015; Weidner *et al.*, 2015), for instance in a recent modelling and experimental study exploring how trade among bacteria can increase bacterial population growth (Tasoff *et al.*, 2015). Studying partner selection and market models in social microbiology is important because, as I pointed out in the introduction to this thesis, biological market models and more in general partner selection models are not mutually exclusively (*e.g.* Wang *et al.*, 2015) with models that explain alignment of fitness between mutualist partners in terms of other factors like

byproducts, partner ‘domestication’, vertical transmission, partner fidelity or spatial structure (Doebeli & Knowlton, 1998; Mueller, 2002; Sachs *et al.*, 2004; Poulsen, 2005; Bshary *et al.*, 2006; Foster & Wenseleers, 2006; Archetti *et al.*, 2011; Verbruggen *et al.*, 2012; Van Dyken *et al.*, 2013; Muller *et al.*, 2014; Kaltenpoth *et al.*, 2014). Rather, it is the relative importance of these factors that needs to be determined. Particularly in systems with mixed-mode transmission, *i.e.* both vertical and horizontal transmission (Bright & Bulgheresi, 2010; Ebert, 2013), multiple factors could contribute to mutualism evolution and stability.

In that sense, even if eventually found to be uncommon, a thorough study of the role of partner selection in microbial mutualism can drive the field forward by excluding one of the potential mechanism for mutualism maintenance as a general factor. I see this teasing apart across different mutualisms of the relative importance of theoretical mechanisms stabilising cooperation, both in contemporary field settings as well as in a deep (historical) context, as perhaps *the* key open question in mutualism research. It will presumably not be answered in the foreseeable future, however. Until then, the most important role for biological market theory in the study of microbial mutualisms (and indeed of all mutualisms), will be to generate new predictions and a theoretical framework applicable to horizontally transferred mutualistic partnerships. This is what I have done in the following chapters of my thesis.

Open question in plant-mycorrhizal mutualisms

Having made the argument that in general biological market models can and should be analysed more widely in microbial systems, I explore in more detail the best studied microbial biological market, the underground exchange of resources between plants and mycorrhizal fungi (Werner & Kiers, 2015a; Chapter 3). Even though there has been substantial theoretical work studying mycorrhizal mutualisms as biological markets (Schwartz & Hoeksema, 1998; Hoeksema & Schwartz, 2003; Kummel & Salant, 2006; Cowden & Peterson, 2009; de Mazancourt & Schwartz, 2010; Grman *et al.*, 2012; Wyatt *et al.*, 2014) and some pioneering experimental work (Bever *et al.*, 2009; Kiers *et al.*, 2011; Fellbaum *et al.*, 2012, 2014; Verbruggen *et al.*, 2012), many questions remain which are discussed in this chapter. A major question that the field should address is what ‘decision rules’ plants (can) use in selecting more beneficial AMF (Bshary & Schäffer, 2002; Sachs *et al.*, 2004; Noë, 2006a), and to what extent this varies across plant species. How (spatially) precise can plants detect individual AMF contributions and reward accordingly? Do plants detect very small quality differences, or is there a minimum threshold? Is expressing a partner preference costly, and how costly? These questions are important because their answers will determine the level of cooperation that evolves. In turn, it is likely that these answers, and the resulting capacity for (precise) partner selection in plants and AMF, depend on the evolutionary history and ecological

conditions of a host species. However, since only a very limited set of model species has been used in this field of research so far (to my knowledge only *Allium vineale*, *Medicago truncatula* and *Plantago lanceolata*) we have little idea to what extent the capacity for partner selection is variable between species, or what factors drive such variation.

In Chapter 3, I highlight the potential for agricultural history as a driver of plant selectiveness and as a tool to study its (short-term) evolution, but one could imagine other potential drivers of variation in plant selectiveness. For instance, it has been suggested that plants that specialised on soils where the benefits of AMF association are low could have reduced interactions with AMF (Lambers & Teste, 2013). It would be interesting to see if these less mycorrhizae-dependent species would also have a smaller need to be selective. Another idea is that root traits, like coarseness or root architecture, could affect the (need for) plant selection of AMF (Grman, 2012), although a phylogenetic comparative analysis revealed no effect of root architecture on plant growth response to AMF (Maherali, 2014).

A third important research direction is studying the ecological relevance and importance of preferential rewarding of co-operators in realistic (field) settings. Plants may be able to select the best-quality AMF in Petri dishes or sterilised greenhouse soil, but how does this work in more realistic conditions? These may include conditions where AMF compete with each other directly as well as through host allocations, conditions where the supply of partners may be limited or conditions where environments fluctuate unpredictably? For instance, if for AMF there is a trade-off between success in direct competition with other AMF and investment in the host plant (Bennett & Bever, 2009; Verbruggen & Kiers, 2010), how does this affect the evolution of cooperation in this system? Presumably, this will depend on the relative balance of factors like AMF abundance, cost of defection versus cooperation for AMF and environmental resources. However, while some of these aspects have now been modelled (Steidinger & Bever, 2014; Wyatt *et al.*, 2014; Bever, 2015), there is only very little work directly addressing (theoretically or experimentally) the potential for a conflict between direct competition among (mutualistic) AMF and plant preferences affecting indirect competition (but see: Engelmoer *et al.*, 2014).

A final major question is if indeed plants (and AMF) always respond as would predicted if they were economic actors to variation in resource abundance. We know that plants select the 'best' AMF in lab settings (Bever *et al.*, 2009; Kiers *et al.*, 2011), but does the strength of such preferential rewarding vary with environmental condition? This is the first question I try to help addressing in the next (empirical) part of my thesis.

CO₂-levels affect relative success of high-quality AMF

It is intuitive to think that changing the relative value of a mutualistically exchanged resource will affect the dynamics of the exchange and the preferences of the trading partners (Pringle, 2015), but it is less intuitive to predict the direction of this change. If we increase the carbon supply to a plant-AMF mutualism, the relative 'price' of a unit of soil nutrients expressed in terms of units carbon will rise (Wyatt *et al.*, 2014; Werner & Kiers, 2015a). Does this mean plants will be more selective in their (potential) preference for high-quality AMF providing lots of these soil nutrients, as predicted by market models (Wyatt *et al.*, 2014)? On the other hand, one could intuitively reason that luxury amounts of carbon release plants from the constraints to allocate it frugally and rather favour them to indiscriminately maximise AMF colonisation. Or alternatively, and returning to the notion of decision rules, perhaps plants decisions are actually not flexible, or at least not in response to fluctuating environmental conditions. That would raise the question if the market metaphor, which assumes flexible allocations to partners based on supply-demand variations, is actually suitable to analyse the plant-AMF mutualism. For instance, if mutualism stability in the plant-AMF mutualism would in fact not rely on preferential rewarding but rather on alternative conceptual models like partner fidelity feedback or byproducts (Sachs *et al.*, 2004), short-term flexible responses to changing environmental conditions would not necessarily be expected.

In Chapter 4, I describe an experiment I performed to determine if varying the carbon availability to plants results in patterns consistent with an effect on preferential rewarding by plants of more cooperative AMF. Despite perhaps being intuitively difficult to predict the direction of the effect, theoretical market models suggest that increasing supplies of carbon selects for higher investment in AMF and more cooperative AMF, because the relative value of soil nutrient increases (Wyatt *et al.*, 2014; Bever, 2015). My first aim with this experiment was to see if the AMF-plant mutualism behaved in accordance with these models. Since AMF are obligate biotrophs and cannot independently acquire carbon (Parniske, 2008), all carbon in the plant-fungal market place is ultimately derived from plant photosynthesis. Therefore, varying the atmospheric CO₂-levels, allows us to experimentally manipulate one of exchanged resources. Additionally, CO₂ is also interesting due to concerns about global change and the associated anthropogenic CO₂-emissions, including potential effects on plant-mycorrhizal mutualisms (Treseder, 2004; Alberton *et al.*, 2005; Drigo *et al.*, 2010, 2013; Fortuna *et al.*, 2012). If indeed environmental fluctuations, like changes in CO₂-levels, can affect plant preferences an important question is how that actually affects the long-term relative success of high and low-quality AMF. After all, we are ultimately interested in notions like partner rewards or sanctions, because these are thought to limit the spread of low-quality partners and stabilise mutualism (Kiers *et al.*, 2011; Oono *et al.*, 2011; Ghoul *et al.*, 2014; Steidinger & Bever, 2014; Bever, 2015). Yet, typically studies analyse plant rewarding dynamics

only over a single generation (Bever *et al.*, 2009; Kiers *et al.*, 2011; Fellbaum *et al.*, 2012; Zheng *et al.*, 2015). Therefore, the second aim with this experiment was to see if there would be any long(er)-term effect of CO₂-levels on relative success of both AMF species in my artificial communities, although I recognise that the three generations I studied represent only a small period in the timeframe of evolution.

My results in Chapter 4 show that, compared to being grown in a monoculture, a low-quality (in terms of effect on host plant growth) AMF species suffered from being grown in competition with a higher-quality AMF, consistent with an effect of plant preferential rewarding. Importantly, this effect was dependent on CO₂-level, with particularly large effects found when comparing depressed CO₂-levels with high and ambient CO₂. Specifically, I found that under depressed CO₂-conditions abundance of both AMF species was reduced, suggesting an absence of plant preferences for more beneficial AMF. In contrast, under ambient and elevated CO₂-levels, abundance of only the lower-quality AMF was reduced, consistent with an effect of preferential rewarding by the host plants. These observations are in line with a scenario where plant preferences are not a fixed response, but rather to some extent vary according to ecological conditions in a manner consistent with stronger selection of AMF under increasing carbon availability. Therefore, the results of this experiment are consistent with the notion that the plant-AMF mutualism can be characterised as an underground biological market and builds on previous experimental (Bever *et al.*, 2009; Kiers *et al.*, 2011; Fellbaum *et al.*, 2014; Zheng *et al.*, 2015) and theoretical results (Hoeksema & Schwartz, 2003; Kummel & Salant, 2006; Cowden & Peterson, 2009).

Previously, long-term CO₂-enrichment has been shown to increase AMF-provided benefits to plants (Gamper *et al.*, 2005). My results suggest that a potential mechanism causing these effects is a more efficient selection of beneficial AMF by plants. This could mean that rising CO₂-levels will strengthen the plant-mycorrhizal mutualism and favour beneficial AMF, potentially mitigating negative effects of global climate change on plant growth. When arguing that my results in Chapter 4 are consistent with market models, I do explicitly note, however, that I only studied patterns: I did not track nutrient flows, or directly measure plant or AMF preferential allocations. However, since AMF obtain carbon from plant photosynthate (Parniske, 2008), relative AMF abundances can be used as valid indication of plant allocations to the fungi (Kiers *et al.*, 2011). For a fuller insight into how plants change preferences in response to shifting environmental conditions, more mechanistic studies can track actual resource flows between plants and AMF (Zheng *et al.*, 2015; Pringle, 2015).

Partner selection may limit long-term spread of low-quality partners

Other than suggesting that plant preferences can vary with environmental context, the second main insight from Chapter 4 is that over multiple generations the relative abundance of a lower-quality AMF in an AMF community is actually successfully reduced. Importantly, my results show this too is likely affected by resource abundance, with both AMF species being rapidly excluded in depressed CO₂-conditions and the high-quality AMF being maintained for longer under ambient and elevated CO₂-levels. This may indicate that preferential rewarding can indeed be effective in reducing the spread of defection in a mutualist population, and that this can be affected by the environment.

Many studies have shown the existence of various forms of partner selection (Adam, 2010; Gubry-Rangin *et al.*, 2010; Oono *et al.*, 2011; Wang *et al.*, 2014), but to my knowledge no studies have experimentally tested the extent to which it is successful in limiting the spread of low-quality partners over multiple generations in mutualisms. While presumably difficult to generalise, because the capacity for partner selection may be difficult to manipulate or shut off in most systems, the strategy I followed in Chapter 4 points towards a potential approach to address this question. By manipulating appropriate environmental conditions that affect (the strength of) partner selection, researchers might create conditions where ‘cheaters’ are more or less likely to spread in experimental populations. The challenge will be to establish that these conditions actually affect partner selection *and* to exclude that changing environmental conditions do not directly differentially affect high-quality and low-quality mutualists.

My present study is limited in scope. I would therefore now argue for additional research into context-dependence of partner selection in the mycorrhizal mutualism, and in mutualisms in general. Chapter 4 represents only an initial attempt to study effects of environmental context on plant preferences: it is limited by the fact that I used only two AMF species that may also differ in other traits. Therefore, I cannot distinguish effects of species identity from effects of AMF partner quality *per se*. To more fully characterise the decision rules that plants (and AMF) follow in rewarding cooperative partners, it will be necessary to study many different species, of varying quality, and across a wider range of environmental conditions including also varying other resources. Such a research programme could further confirm, or exclude, the applicability of market models to the AMF-plant mutualism. Data from more extensive studies across different (combinations of) plants and mycorrhizal species and environmental conditions could also allow us to draw the curves that determine how partner rewards depend on relative price of a mutualistically provided benefit, which potentially has important potential implications for the level of cooperation that evolves in mutualisms (Bshary & Schaffer,

2002; Sachs *et al.*, 2004; Noë, 2006a). This is an interesting question to address not only within the plant-mycorrhizal mutualism, but also across different mutualisms.

A second layer of complexity is that in Chapter 4, and in this thesis in general, I have typically referred to mutualistic partners as ‘high-quality’ or ‘cooperative’ versus ‘low-quality’ or ‘less-cooperative’. However, considered from a biological market perspective, these descriptions cannot be intrinsic to any mutualist, but are inherently context-dependent. For instance, a plant acting as an economic actor would reward an AMF species efficient at providing nitrogen in nitrogen-depleted soil, while in other soils a phosphorus-providing AMF may be a ‘high-quality’ mutualist (Werner & Kiers, 2015a). Thus, it is not just the preferential rewarding of one partner which may be affected by changing environmental conditions, but potentially also the identity of the partner which is of ‘high-quality’.

Competition between AMF can limit the scope for plant partner selection

My next chapter addresses the existence of priority effects among AMF species colonising plant roots. In ecology, priority effects refer to the situation where a species arriving early in a habitat can successfully dominate that habitat in competition with other species, even if these would have been superior had they arrived concurrently (Chase, 2010; Fukami & Nakajima, 2011). Initially, this chapter might not seem to clearly fit within the general biological market framework of my thesis. However, even though biological market theory typically does not consider direct competition among mutualists but rather focus on indirect (outbidding) competition over partner benefits (Noë & Hammerstein, 1994, 1995), direct competition could in principle affect mutualism strength and the stability of cooperation (See also Chapter 3). Suppose a species of AMF is superior in direct (*i.e.* not through plant mediation) competition with other species, for instance through priority effects. This could cause other AMF to disappear from the population or become substantially reduced, effectively limiting the set of AMF available for plants to select a mutualistic partner from (Werner & Kiers, 2015a). Thus, even though I have emphasised in this thesis the role of partner selection driving biological market dynamics, direct interactions between mutualists in the same ‘trading class’ (like various AMF species) will also affect these dynamics. There is a potential analogy here between biological markets and human economies. In economics, the ‘first-mover advantage’ refers to a situation where due to technological leadership, consumer switching costs and acquisition of scarce assets, firms which arrive first in a market are able to outcompete later arrivals, even if these would otherwise be more efficient (Spence, 1981; Lieberman & Montgomery, 1988, 1998). Therefore, if they exist within a mutualistic exchange, priority effects could be seen as the biological analogue of first-mover advantages.

In chapter 5, I describe an experiment performed to test if priority effects exist among AMF colonising a young seedling. The results I present reveal that the first arriving AMF species dominates root colonisation in young plants, regardless of the identity of the species (Werner & Kiers, 2015b). Since I used the same pair of AMF species where in previous research *Medicago truncatula* plants were found to preferentially allocate resources to the more cooperative species (Kiers *et al.*, 2011; Fellbaum *et al.*, 2014; Chapter 4), this suggests that arrival order might counteract the host plant's ability to effectively select AMF partners, because even the low-quality AMF species can exclude a later arriving invading AMF. Therefore, the results in this chapter call for caution in extrapolating from artificial, lab-based choice experiments to field settings. Partner choice can be counteracted by other factors like direct (antagonistic) competition and priority effects, and ecologically realistic experiments are required to determine its importance in the field and its effectiveness in checking defecting mutualists.

For instance, if a less-cooperative mutualist is strongly superior in direct competition, that could represent a viable strategy even while potential mutualistic partners favour more cooperative individuals. There is some support for a trade-off in AMF between competitive superiority and benefits for host plants (Avio *et al.*, 2006; Bennett & Bever, 2009; Chagnon *et al.*, 2013), and this could help explain why despite preferential rewards for cooperative AMF, low-quality and even parasitic AMF still persist. Potentially, specialising on competitive superiority versus on being cooperative to the plant represent co-occurring AMF strategies: my results suggest that specialising on rapid colonisation of young seedlings could be a component of this competitive superiority, as has also been suggested in ectomycorrhizal fungi (Lilleskov & Bruns, 2003). I studied only a relatively short time period (twelve weeks): in the long run, initial priority effects may still be overcome. Therefore, a potential AMF strategy of rapid colonisation may be particularly successful in habitats with a high host plant turnover, such as areas with a lot of disturbance, while long-term stable habitats might offer more opportunities for plant selection of AMF to become effective. If true, this would predict that more stable ecosystems, dominated by longer-lived plants, will have more cooperative AMF.

My research in Chapter 5 raises various other new questions. For instance, what is the relationship between priority effects among AMF within roots and within the soil? While I have only quantified intraradical abundances in my work, AMF are also interacting extraradically in the rhizosphere. This extraradical competition may mirror intraradical colonisation dynamics, especially if the resident can monopolise access to plant carbon and thereby produce more hyphae than the invader. Under this scenario, feedback loops could arise further strengthening the advantages of the resident AMF. Using large-scale genotyping (Dumbrell *et al.*, 2010, 2011) and AMF interaction network analysis (Chagnon *et al.*, 2012), researchers have suggested that early random colonisation of

certain AMF species can drive a positive feedback loop, where the pioneer species benefits most from the plant carbon supply and therefore has an increased chance of further colonising local plants compared to species arriving later. My experimental design with a single plant per pot might have been particularly favourable for such feedback loops to arise because after initial carbon supplies in spores have been exhausted, the single plant becomes the only source of carbon. This could restrict the opportunities for repeated invasion compared to a more natural setting in which most AMF will be connected to other plants through their mycelial network (Giovannetti *et al.*, 2004) and would thus have access to alternative carbon sources. Future work is needed to determine how common mycelium networks affect the strength of priority effects.

Another relevant question was highlighted in a commentary to my work presented in Chapter 5 (Johnson, 2015). In contrast to most other examples of priority effects (*e.g.* Wilbur & Alford, 1985; Weslien *et al.*, 2011; Dickie *et al.*, 2012), roots, the habitat that AMF colonise, are another biotic ‘sub-system’, potentially also subject to priority effects (Johnson, 2015). This means there is the potential for priority effects in both systems to interact, potentially generating much more complicated dynamics than I have highlighted here. This idea also touches upon the potential impact our results have on effort to use AMF inocula to stimulate crop yield in an agricultural setting (Verbruggen *et al.*, 2013). My results could imply that in (agricultural) systems with well-established resident AMF networks in the soil, priority effects among AMF could limit the effectiveness of additional inocula. This will typically be systems where plants are well established too (and thus experience potential advantages from priority effects), while systems with recent disruption (such as through tilling) could be more susceptible to inoculum addition.

Lastly, I want to highlight a potential conflicting result within my PhD thesis. While in Chapter 4, we found that our ‘low-quality’ AMF (*Glomus aggregatum*) had a small, but statistically significant, different effect on plant growth than the higher-quality AMF species (*Rhizophagus irregularis*), in Chapter 5 we did not find a difference in plant growth response between these two AMF species (Werner & Kiers, 2015b). Because of this, I have not directly discussed my results in Chapter 5 in the light of AMF cooperation level or plant preferential rewarding, although the general point that priority effects exist and can be a potential limitation to plant partner selection remains valid even if in the experimental conditions used there were no actual quality differences between both AMF. Yet, this does emphasise that to truly measure the extent to which priority effects counteract plant preferences, and reverse, one would ideally need to study various AMF with varying quality differences as well as various head starts. Secondly, the observation that (relatively minor) variations in the conditions between both experiments can influence the effects of AMF on plant growth highlight that, as I have argued previously

in this discussion, notions of (AMF) partner quality are not intrinsic to the species studied but only have meaning relative to the particular environmental conditions.

Division of labour among symbionts can be favoured by active host selection

An open, and rarely asked, question in mutualism research is to what extent organisms exert control, direct or indirect, over *which* actual benefits potential mutualists offer. In contrast, studies typically take the mutualistically provided service or resource as a given, and study the level of investment in the mutualistic partner, *i.e.* study how much of a given benefit a mutualist provides (*e.g.* West *et al.*, 2002a; Wyatt *et al.*, 2014). Yet, *what* is provided is potentially as much an evolving variable, subject to natural selection, as the *level* of cooperation. For instance, many hosts obtain multiple services or resources symbiotically: plants are pollinated by insects, receive different soil nutrients from a range of different (microbial) mutualists, and can be protected by ants (Ness, 2006; Parniske, 2008; Rodriguez *et al.*, 2009; Oldroyd, 2013). Interpreting this in terms of (biological) market metaphors, this can be seen as the equivalent in human economies of some firms specialising on a single or a few products, while others offer a diversified range of products and services (Capon *et al.*, 1988; Borghesi *et al.*, 2007). In biology, however, we have little idea why some mutualists, like AMF, provide a range of services (Augé, 2001; Parniske, 2008; Veresoglou & Rillig, 2012), while other hosts have a division of labour among specialised symbionts (Bshary & Grutter, 2002b; Davis *et al.*, 2014). Therefore, I have created the first formal game-theory model addressing this question, which is discussed in Chapter 6. This general model is also applicable to most plant underground mutualisms and indeed to any mutualism where a host has the potential to interact with multiple types of symbionts.

The model reveals that there are two crucial factors in determining if a symbiotic division of labour evolves. First, the trade-off symbionts face between the various (in the model two) mutualistic services is key: if they are complementarily transferred to the host (*e.g.* when one service is a by-product of the second) a division of labour is never favoured by evolution because a multifunctional symbiont can always transfer more resources (and receive a higher fitness in return) at relatively small cost. In contrast, if the mutualistic services considered are antagonistic (symbionts need to invest in the one at the expense of the other), division of labour can be favoured but only if there is a high number of symbionts per host to ensure that the chance of the host lacking a crucial resource is small. However, I also establish that high symbiont numbers per host can cause the breakdown of symbiont cooperation with the host. The reason for this is that, from the perspective of the symbionts, the host is a collective resource and high symbiont numbers dilute the individual symbiont's contribution to this resource, resulting in a 'Tragedy-of-the-Commons' (Hardin, 1968; Rankin *et al.*, 2007).

We find that both host preferential rewarding of more cooperative symbionts, the strategy predicted to evolve in biological markets (West *et al.*, 2002b,a; Wyatt *et al.*, 2014), and active selection by the host of a diverse symbiont community (in terms of mutualistic services offered), can stabilise symbiont cooperation and can enable the high symbiont numbers necessary for a division of labour to evolve. This shows that partner selection can not only affect the evolution of the level of cooperation of a mutualist *per se*, but also the nature of that cooperation. While the model I present in Chapter 6 is simple, and does not capture all sorts of potentially relevant aspects of typical mutualisms (such as more than two required resources, potential for benefits to be obtained non-symbiotically or decreasing marginal benefits of additional resources), to my knowledge, this is a possible role of partner selection which has not so far been explored in the literature.

A cryptic and ancient precursor innovation is crucial in nodulation evolution

So far, my research has taken the prior existence of a mutualism as a given and focused on mechanisms stabilising the partnerships and on (environmental) factors affecting these stabilisation mechanisms. However, how do mutualisms arise? In Chapter 2, I argued that a key open question is how mutualistic exchanges evolve from non-cooperative interactions. One way to help address this is take a historical perspective and study the evolutionary history of the complex adaptations and organs that (plant) hosts have evolved to accommodate their (microbial) mutualists. In the last two chapters of my thesis, I take this historical perspective and use phylogenetic comparative methods to study origins (Chapter 7) as well as subsequent evolutionary persistence (Chapter 8) of plant-microbe mutualisms.

I here switch to studying the mutualism of plant and N₂-fixing bacteria. These mutualisms are characterised by specialised, and highly complex, root nodules that house the bacteria (Masson-Boivin *et al.*, 2009; Oldroyd, 2013). In contrast to in plant-mycorrhizal mutualisms, plants are more in control in this system, as they can detect and senesce root nodules not providing (sufficient) nitrogen, while the bacteria are physically enclosed in nodules and (until the nodules senesce) cannot form a new partnership (Kiers *et al.*, 2003, 2006; Simms *et al.*, 2006; Gubry-Rangin *et al.*, 2010). These plant sanctions are thought to help stabilise cooperation by the bacterial partners, stimulating bacterial investment in the costly process of N₂-fixation (West *et al.*, 2002b; Denison & Kiers, 2004; Kiers & Denison, 2008). One idea is that for sanctions to be effective in mutualisms, symbionts cannot be mixed. Rather, hosts need to have a way of spatially separating or compartmentalising symbionts in order to be able to evaluate (individual) symbiont contributions (Frank, 1996; Verbruggen *et al.*, 2012). For instance, bobtail squids house bioluminescent bacteria in physically separate ducts and carefully monitor bacterial access to these crypts (Nyholm & McFall-Ngai, 2004; McFall-Ngai *et al.*,

2012; Soto *et al.*, 2012). Similarly, insects separate their endosymbionts in specialised compartments called bacteriocytes (Matsuura *et al.*, 2015) where they can be controlled (Nyholm & Graf, 2012; Koga *et al.*, 2013), while stinkbugs have a specialised intestinal organ for symbiont sorting (Ohbayashi *et al.*, 2015). In symbiotically N₂-fixing plants, nodules allow for this separation: there is a complex pathway to initiate their formation and ensure that only a single (or at most a few) clones enter a given nodule (Masson-Boivin *et al.*, 2009; Oldroyd *et al.*, 2011; Oldroyd, 2013). A key question that remains however, is when and how the nodules that allow for and mediate effective partner control (and sanctions) first evolved.

In Chapter 7, I present the first evolutionary reconstruction of nodulation evolution across the angiosperms, going back over two hundred million years in time, and identify the key events that enable nodule evolution. After compiling the world's largest database of symbiotically N₂-fixing angiosperms (Werner *et al.*, 2014b), and using the most comprehensive phylogeny of global plants species (Zanne *et al.*, 2013, 2014), I used a novel quantitative framework to reconstruct the evolution of nodulation (Beaulieu *et al.*, 2013; O'Meara & Beaulieu, 2014). The analyses reveal an ancient, cryptic 'precursor' to symbiotic N₂-fixation. The precursor is an evolutionary innovation which did not itself produce the mutualistic phenotypes in the plant species possessing it, but which enabled its evolution. The actual origins of nodulation in angiosperms evolved dozens of millions years later. This shows that there is likely to be a 'deep homology' across all instances of nodulation in angiosperms, with an ancient and conserved genetic innovation enabling mutualistic N₂-fixation (Shubin *et al.*, 2009). This is interesting because it suggests that anciently conserved regulatory circuits cannot only underlie complex adaptations like eyes, limbs or leaves (Shubin *et al.*, 1997, 2009; Pires & Dolan, 2012), but also adaptations that function to mediate interactions with, presumably co-evolving, (mutualistic) partners.

While I only detected the phylogenetic signature of the symbiotic N₂-fixation precursor, not its identity, the notion of a very rare but crucial innovation can help explain what I think is one of the most interesting puzzles this mutualism represents: why is it so rare? In contrast to arbuscular mycorrhizal mutualisms, which are distributed widely across almost all major clades of land plants (Parniske, 2008; Humphreys *et al.*, 2010; Delaux *et al.*, 2013), nodulation is rare (taxonomically, not in terms of number of species) and found in only four orders (Doyle, 2011). Yet, we know that nitrogen is one of the key factors limiting global plant growth (Reich *et al.*, 2006; Lambers *et al.*, 2008; Batterman *et al.*, 2013a), and in many cases plants greatly benefit from symbiotically fixed nitrogen (Thrall *et al.*, 2000; West *et al.*, 2005; Batterman *et al.*, 2013a). Consequentially, one of the most remarkable puzzles about the plant N₂-fixing mutualisms is that it is not more widely found. My reconstructions reveal that precursor evolution was a very rare

event, suggesting that the precursor represents a complex genetic innovation, unlikely to evolve multiple times. The key question will now be to find the genetic correlate of the precursor, also because this might hold the key towards the long-standing aim of engineering N₂-fixing crops, like cereals (Delaux *et al.*, 2015b).

Comparative methods can help elucidate mutualism evolution

There is an interesting parallel between the N₂-fixation precursor I report and the evolution of other complex traits, which can also require ‘pre-adaptations’ or ‘predispositions’ that do not directly cause the trait studied, but enable its evolution (Rajakumar *et al.*, 2012; Blount *et al.*, 2012; Marazzi *et al.*, 2012; Barve & Wagner, 2013; Quandt *et al.*, 2014). To my knowledge, this chapter is both the first to show the large-scale phylogenetic signal of a pre-adaptation using ancestral trait reconstruction (rather than inferring it from molecular or genetic evidence, usually in a few species) as well as one of the first to show pre-adaptations for mutualistic traits (for another example, see: Marazzi *et al.*, 2012). Since, another study has also used comparative approaches across large numbers of species to show that in fact the algal ancestor of land plants was already pre-adapted for fungal symbioses, which later evolved into fully functional mycorrhizal symbioses (Delaux *et al.*, 2015a). These approaches are also potentially relevant for the debate in evolutionary biology about the importance of convergent or repeated evolution versus historical contingencies (Ord & Summers, 2015), because pre-adaptations or precursors could be one of the factors driving repeated evolution of the same phenotypes in related taxa (Rajakumar *et al.*, 2012), in addition to similar selection pressures. I did not address to what extent similar selection pressures were responsible for the about eight origins of symbiotic N₂-fixation we found, but my results reveal that the repeated evolution of the nodulation phenotype was constrained by the necessary (genetic) background of the single evolutionary precursor.

While phylogenetic comparative methods cannot directly resolve mechanistic questions, Chapter 7 demonstrates how, in combination with large databases and phylogenies, they can be exceptionally powerful to identify previously hidden, crucial steps in the evolution of a trait many millions of years ago. Uniquely, and in contrast to verbal accounts of evolutionary history, this allows also for quantification of the certainty associated with a particularly evolutionarily scenarios. The advantage of this is that it can help us find interesting species to study, for instance highly likely current precursor species, which may be key to elucidate the mechanistic and molecular basis of complex traits (Werner *et al.*, 2014a; Delaux *et al.*, 2015b). Thus, Chapter 7 might not tell us where to look *genomically* (*i.e.* where in the genome), but it can help us tell where to look *taxonomically* (*i.e.* in what species) in for instance a comparative genomics approach. More in general, I would argue that advanced quantitative phylogenetic reconstructions, like the ones used here, represent very promising avenues in mutualism evolution and indeed in

evolutionary biology broadly. The reason for this is that because of the quantitative nature of these reconstructions, we can now correlate evolutionary reconstructions for a (mutualistic) trait like nodulation to potentially explanatory factors driving its evolution. Examples could be factors like historical environmental conditions (using paleontological records), other organism traits and the evolution of partner species.

Potentially, this will help us address the question I raised earlier in this Synthesis: how important are various theoretical mechanisms of mutualism stability, not just in single species pairs and in a limited experimental context, but across (thousands of) species and over deep time? While causality will always remain difficult, even impossible, to disentangle in a research programme along these lines, phylogenetic correlations can exclude potential explanations. For instance, if a large-scale phylogenetic comparison analysing the capacity for partner selection and the stability of mutualisms fails to find a positive relationship between these, this suggest that other factors than rewarding or sanctions may have been more important in those systems in driving stable cooperation. Consequentially, I see the combination of large-scale phylogenetic analyses across mutualisms, in combination with experimental (and field) work tracking resource flows to quantify mutualistic benefits and with experimental evolution approaches confirming or excluding causality, as the way forward towards a general and comprehensive understanding of mutualism evolution across the tree of life.

Persistence and breakdown of the legume-rhizobia mutualism

In the final chapter of my thesis, I have attempted a contribution to this research programme (Werner *et al.*, 2015). My analyses in Chapter 7 revealed the existence of a previously unknown class of plant species where the mutualism with N₂-fixing bacteria is extraordinarily stable, and is extremely unlikely to be lost even over evolutionary times (Werner *et al.*, 2014a). Why is mutualism stable in some lineages, but easily lost in some others? Most studies in the field of mutualism research, for understandable reasons, study those systems where we observe (stable) cooperation. But in many ways, even more interesting are those interactions where the partnership failed and broke down (Sachs & Simms, 2006). I expect we can learn more about what makes a stable mutualisms by trying to determine what discriminates those lineages where cooperation failed from those where it is a continuing success and persists. But how do we detect failed mutualisms? While it may be impossible to find those systems where mutualism broke down and species went extinct as a result (Sachs & Simms, 2006), we can find non-mutualistic clades nested within mutualists (*e.g.* Nelsen & Gargas, 2007; Sachs *et al.*, 2011b), sometimes even in the case of very intimate symbioses such as organisms losing mitochondria or chloroplasts (Merckx & Freudenstein, 2010; Hjort *et al.*, 2010). Additionally, my results in Chapter 7 and 8 show how advanced phylogenetic comparative methods can be used to detect clades that, while they still

have the mutualism, have variable mutualism loss rates and help find lineages with stable and more unstable mutualisms.

Using this reasoning to my advantage and reanalysing my results from Chapter 7, I show that, in contrast to my expectations (Lambers *et al.*, 2008; Lambers & Teste, 2013), soil nutrients do not predict stability of the plant-rhizobial mutualism in legumes (Chapter 8). However, we find that plant traits, specifically high leaf nutrient concentrations, are associated with mutualism stability. While I again emphasise that causality cannot be inferred from this observation, and while ideally many more factors would be incorporated in a large-scale analysis of potential explanatory variables, the second result was expected. The reason for this is that, we would expect that the benefits obtained from a mutualistic lifestyle would be a key factor determining if loss of the symbiosis is favoured (Noë & Hammerstein, 1994, 1995; Wyatt *et al.*, 2014). High nutrient concentrations, and thus a highly nutrient demanding lifestyle, would increase this benefit from mutualism for plants. My results in Chapter 8 suggest that this causes legume plants to become more likely to be 'locked' into a symbiotic lifestyle and reveal how long-term mutualism persistence and stability is affected by a complex interplay between host traits and environmental factors

This suggest that an interesting open question here is to consider other factors that mediate the relative costs and benefits of engaging in the mutualism, such as solar radiation intensity, temperature or CO₂-levels. Similarly to in the arbuscular mycorrhizal mutualism (Wyatt *et al.*, 2014 Chapter 4), from a cost/benefit perspective one would expect higher availability of carbon to cause soil nutrients to be relatively more limiting and costly. This would favour stability of the mutualism providing these nutrients because of increased benefits from mutualism, and possibly even dependence of the host plant on the symbionts. However, when analysing the effect of temperature on likelihood of being a highly persistent mutualists, I found the reverse effect, with lower temperatures suggesting higher mutualism stability. Similarly, and in contrast to our expectations, we did not find an effect of soil nutrient content on mutualism stability. I currently do not have a clear explanation for these surprising effects. A way forward may be to collect much more detailed and spatially precise soil data than our current analysis at a large spatial scale (1/2 a degree, *i.e.* a little over 50 kilometres in the North-South direction). However, collecting this, and in general collecting trait or environmental data across thousands of species will be challenging, all the more so because ideally analyses like these should also take into account intraspecific variability (Ives *et al.*, 2007; Garamszegi & Møller, 2010).

A second class of traits that would be very interesting to consider when describing mutualism persistence would be traits that help control partner cheating, for instance

the capacity for effective sanctioning. Here too, a limitation is that large-scale datasets for this are not available. In general, other than considering (i) traits that influence partner selection and (ii) factors that influence mutualism cost/benefits a third approach would be to study evolution on the bacterial side of the interactions. In principle, mutualism stability in the legume-rhizobia mutualism can depend on bacterial traits too, and I have not considered these. For all these reasons, my work in Chapter 8 represents only the beginning of what can be done to study mutualism persistence and identify the factors distinguishing stable from unstable cooperation.

Conclusion

This thesis contributes to a wide range of open questions in mutualism evolution, from the applicability of an important class of mutualism models to microbes and specifically AMF (Chapters 2 & 3), experimentally disentangling ecological factors that can affect underground mutualism stabilisation (Chapters 4 & 6) to theoretical analysis of the nature of the (underground) mutualistic exchange that is expected to evolve (Chapter 6) and to elucidating the historical origins, predispositions and persistence of these cooperative interactions (Chapter 7 & 8). I have used plant-microbial mutualisms to draw a general main conclusions regarding the evolution of cooperative behaviours between horizontally-transmitted symbiont species. Most important among these are: (i) biological market theory can be a useful framework across a wider-range of organisms than it has traditionally been used in, including microbes (Chapters 2, 3 & 4), (ii) mutualisms can show flexible partner preferences which may be modified according to environmental conditions and which can affect mutualist's long-term success (Chapter 4), (iii) success on a mutualistic exchange is not only driven by partner preferences but can also be affected by direct competition among potential partners (Chapter 5), (iv) active partner preferences not only affect level of cooperation but can also shape other aspects of a mutualism like the level of mutualist specialisation (Chapter 6), (v) the origins of mutualistic cooperation can be predated by complicated evolutionary trajectories including ancient, cryptic 'precursors' (Chapter 7) and (vi) long-term mutualism persistence and stability is affected by a complex interplay between host traits, symbiont traits and environmental factors (Chapter 8).