The Reinhart-Heinrich Doctoral Thesis Award 2018

The winner of the ESMTB Reinhart-Heinrich Doctoral Thesis Award for 2018 is Computer Science graduate Daniel Nichol, for his thesis **Understanding drug resistance through computational models of the genotype-phenotype mapping**. The award will be formally given in a ceremony at the ECMTB conference in Heidelberg in 2020. The awarding committee made the following statement: *Daniel Nichols thesis is likely to become a landmark in an issue of critical significance. It studies and proposes strategies to deal with drug resistance in the treatment of cancers and infectious diseases taking into account the genotype-phenotype mapping. New mathematical models are used to better understand the underlying mapping mechanisms, integrating genetics, environmental signals and stochasticity. The thesis gave rise to three papers already published in highly prestigious journals and to a few more papers under publication.* Dr Nichol completed his DPhil at University College Oxford. His supervisors were Peter Jeavons (Oxford, Computer Science) and Alexander R. A. Anderson (Moffitt, Integrated Mathematical Oncology). We congratulate him for his excellent and exciting work! First, he presents himself, and then follows an extended abstract of the thesis.

Daniel Nichol Personal statement

To predict this evolution we must first understand how genetic mutations are manifest as phenotypic change, and how these phenotypes are selected by druginduced environmental change. Studying this 'genotype-phenotype mapping' lies at the heart of my DPhil studies and my later modelling work as a postdoctoral training fellow in the Centre for Evolution and Cancer at the Institute of Cancer Research, London. I believe that it is only through an interdisciplinary approach combining modelling, large scale genomics, and experimental evolution that the evolution of drug resistance can be fully understood. Ultimately, if we can better predict the evolution and ecology of cancers, then we may be able to design better therapies by exploiting evolutionary trade-offs. The research I have undertaken in my DPhil was the first step towards this goal.

Thesis summary: 'Understanding drug resistance through computational models of the genotype-phenotype mapping' by Daniel Nichol

The emergence of drug resistance is ultimately driven by Darwinian evolution. These evolutionary dynamics are inherently two-tiered, with mutational processes at the genetic scale inducing variation in cellular phenotypes that are subject to natural selection. If we are to predict or reverse evolution, as we must to determine effective treatments for drug-resistant infections and cancers, then we must first understand the relationship between genetic and phenotypic change. This relationship, known as the genotype-phenotype (GP) mapping, is governed by a complex cascade of potentially stochastic molecular interactions that integrate genetic, epigenetic and environmental factors to determine cellular phenotype. In this thesis, we introduce mathematical models of the GP-mapping to explore how its structure influences the evolution of drug resistance, and how it determines the efficacy of novel therapeutic strategies such as drug holidays or adaptive therapy.

We begin by providing a comprehensive review of previous abstract modelling of the GP-mapping. Second, we demonstrate that models of GP-mappings can improve evolutionary predictions. Specifically, we show that through careful selection of drug sequences evolution could be 'steered' such that a highly drug resistant population does not emerge. Third, we introduce a novel model for

the GP-mapping wherein phenotypes arise stochastically. Through this model we explore the evolutionary dynamics of ecological 'bet-hedging', a common driver of drug resistance. We find that the structure of the GP-mapping can slow the evolutionary loss of this trait, preserving the survival mechanism when harsh environments occur very infrequently. Thus, the capacity to steer the evolutionary loss of bet-hedging is dependent on the structure of the GP-mapping. We next extend the bet-hedging model to account for epigenetic inheritance and the potential for phenotypic memory. Critically, we find that genetics, epigenetics and the GP-mapping interact non-additively to determine organismal fitness, indicating that evolution likely cannot be predicted without accounting for each of these biological processes. Finally, we explore how properties of the GP-mapping are manifested at the population scale and suggest experimental approaches to identify bet-hedging through population scale assays. The implications of our results for the treatment of drug-resistant diseases are explored throughout. Here, we provide an extended summary of the contents of the full thesis.

Introduction and Overview

The evolution of drug resistance represents a growing health crisis in the treatment of infectious diseases that necessitates immediate intervention. Similarly in the treatment of cancers, evolved secondary drug resistance is common, ultimately driving mortality. In both cases the same underlying evolutionary dynamics are at work. Therapeutic intervention induces strong selection pressures on a phenotypically heterogeneous population, driving the emergence of a resistant population and ultimately the relapse of refractory disease [1]. Evolutionary therapy aims to mitigate this process through timed interventions with alternative drugs, or breaks from drugs, informed by mathematical modelling. Critically, the success of this approach is dependent on our ability to predict evolution, and in particular, to predict when and how resistant phenotypes will emerge [2].

Phenotypic heterogeneity is often thought of as arising from rare, random mutations. However, isogenic populations can also generate phenotypically heterogeneous populations, either through stochastic cell-fate determination (bet-hedging), or through modulation driven by environmental cues (phenotypic plasticity). The relationship between genetics, environment and phenotype is encapsulated in the classical concept of the genotype-phenotype (GP) mapping which has been explored in numerous biological subfields [3]. Here, we propose that understanding the GP-mappings that underpin drug resistance will lead to better evolutionary predictions, and thus better therapy. To this end, we set out to derive GP-mapping models of the drivers of drug resistance, to explore their implication for evolutionary therapy, and ultimately derive conditions under which the drivers of adaptive phenotypic heterogeneity can be identified and potentially mitigated.

Aims: The aims of this thesis are three fold: First, to unify previous mathematical modelling of the GP-mapping in order to identify those sources of phenotypic heterogeneity that have been under-studied with respect to the evolution of drug resistance; second, to implement new models of the GP-mapping that account for these sources of phenotypic heterogeneity; and third, to demonstrate that coupling models of the GP-mapping with previous models of population dynamics can improve evolutionary predictions. This third aim is tackled three times with increasingly complex GP-mappings, each time starting from an abstraction of evolution and building theory towards experimentally testable predictions.

Experimental validation: This work was completed in close collaboration with experimentalists at the Moffitt Cancer Center (Tampa, FL), the Cleveland Clinic (Cleveland, OH), and Case Western Reverse University (CWRU) (Cleveland, OH). These collaborations helped to ensure that throughout the thesis we always worked toward generating experimentally verifiable hypotheses. Owing to time constraints, experimental results were not complete before these thesis was submitted but have since validated a number of theoretical predictions and been published. Where appropriate

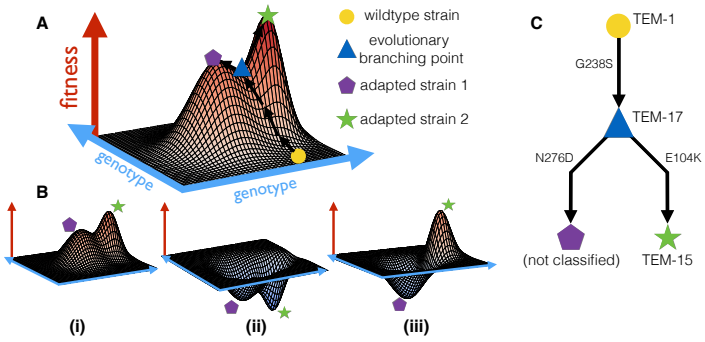

Figure 1: **Collateral sensitivity need not be repeatable. A)** Evolution from a wild–type genotype (yellow circle) need not be repeatable owing to the potential of branching points in the fitness landscape (blue triangle). **B)** Patterns of collateral response that can arise from evolutionary divergence: **i)** guaranteed cross–resistance, **ii)** guaranteed collateral sensitivity, **iii)** collateral response determined by which evolutionary trajectory occurs. **C)** A model predicted evolutionary branching point in the TEM gene under the antibiotic cefotaxime [10].

these validation experiments are highlighted below.

GP-mappings and Phenotypic Heterogeneity

The genotype-phenotype (GP) mapping is central to successful predictions of evolution. With the advent of single-cell sequencing it becoming clear that whilst mutations often drive phenotypic diversity, monoclonal cells exhibit phenotypic heterogeneity. Deconvolving the drivers of this heterogeneity is a difficult task, owing to the overwhelming complexity of intra-cellular signalling cascades. Mathematical models provide a means to understand this complex process, just as modelling has helped elucidate other complex ecological and evolutionary systems. Models of the GP-mapping arise many biological fields, often with recurrent properties, permitting us to gain intuition for one system from the study of another. For example, in attempting to understand drug resistance, studies of RNA secondary structure provide insight into the role of neutral mutations and studies in plant biology provide mathematical formulations of the 'reaction norm' governing phenotypic plasticity. The mathematical basis of these results is independent of the specific domains in which they were derived. We present a unification of GP-mapping results under a common framework through which evolutionary questions can be interpreted beyond the fields in which they originally arose. Our survey highlights aspects of the GP-mapping that are under-studied evolutionary medicine; namely bet-hedging which forms the basis of the much of this thesis.

Predicting Evolution from GP-mappings

The central tenet of our work is that models GP-mappings can improve evolutionary prediction. To demonstrate this fact we began by exploring the simplest GP-mapping - the fitness landscape. By consideration evolution under the constraints of a fitness landscape, we showed that drug ordering determines the degree of resistance that arises. We utilised empirically-derived fitness landscapes for *E. coli* to provide evidence that rational drug sequencing can partially mitigate resistance, whereas arbitrary sequencing may promote it [4]. Absent clinical guidelines regarding drug sequencing we argue that current practice may promote drug resistance just as irresponsible dosing can.

Our modelling suggests that evolution may be predictable but not necessarily repeatable (Figure 1). Evolution need not arise in the same way in independent instances, but evo-

lutionary trajectories can be sufficiently constrained for predictions to be made. An immediate consequence of this observation is that predictions of collateral sensitivity made using low-replicate-number evolutionary experiments may have overstated therapeutic benefit; predicting a collaterally sensitive response where cross resistance can occur. As such, we introduce the concept of *collateral sensitivity likelihoods* as a metric to account for stochasticity in the evolution of drug resistance. Following the submission of the thesis, this non-repeatability was confirmed through high-replicate experimental evolution coupled with targeted whole-genome sequencing [5]. These results have motivated changes in our own experimental evolution protocols undertaken in collaboration with microbiologists at CWRU, Ohio. We anticipate they will also impact how others will undertake experimental evolution studies in future.

Non-genetic Heterogeneity and the Evolution of Bet-Hedging

Resistant phenotypes can arise stochastically within an isogenic population. A clinically important example is the emergence of drug-tolerant bacterial 'persister cells' that allow bacterial populations to survive antibiotic therapy before outgrowing to re-establish a pathogenic population [6]. In our review of GP-mappings we identified that this phenomenon is under-represented in theoretical models. Inspired by the milieu of intracellular chemical reactions that govern cellular phenotypes, we implemented a new model of bet-hedging using stochastic chemical reaction network models previously studied in the context of biological computing [7] (Figure 3). Specifically, we modelled phenotypes as determined by the stochastic simulation of a bistable switch with genetically determined initial conditions [8].

Mathematical studies have demonstrated that bet-hedging can be selectively advantageous in stochastically fluctuating environments but can incur a fitness cost in fixed environments. This represents an apparent paradox of bet-hedging: how can bet-hedging persist over long time-scales in environments where natural selection acts to remove it? We performed invasion analysis to demonstrate that the structure of the GP-mapping itself can act to slow the evolutionary loss of bet-hedging by inducing diminishing fitness benefit in successive mutations towards a one-phenotype strategy. This result is not solely an interesting evolutionary result, but has profound impact for evolutionary therapies targeting diseases with bet-hedging-driven resistance mechanisms. Through agent based modelling of populations subject to drug holidays, we demonstrated that where evolutionary loss is fast, treatment holidays may induce increased susceptibility to drugs targeting proliferative cells, whereas where evolutionary loss is slow, this strategy is likely to fail. Thus, we argue that understanding of the precise molecular drivers of bet-hedging may hold the key to predicting and reversing some important forms of drug resistance.

Non-genetic Inheritance and Evolution of Phenotypic Memory

Previous models of bet-hedging assume that offspring are assigned a phenotype with a fixed probability. Often, this probability is taken as dependent on the parent cell phenotype, which represents a type of stochastic inheritance or 'memory'. We set out to explore how this memory may arise and whether it impacts the evolutionary fate of a bet-hedging population. We first extended the bistable switch model by introducing the epigenetic inheritance of the intracellular factors that comprise the switch. These inherited molecules interact with the genetically-determined initial 'burst' of expression to bias the switching behaviour (Figure 3). Once the switching has occurred, the molecules are modelled as subject to decay throughout the life of the cell, in turn determining the number inherited epigenetically at the next reproduction event. This mode of epigenetic inheritance is observed in human cell lines where the relative abundance of epigenetically inherited stress protein (p53) and mitogen (CCND1) determine cell-cycle arrest or proliferation [9].

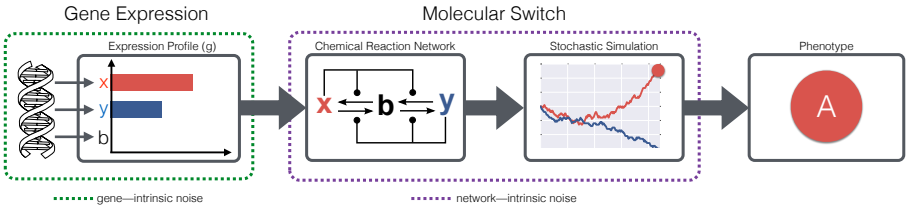

Figure 2: **The chemical reaction network model for bet-hedging.** Gene expression profiles are assumed to be fixed for each genotype and serve as initial conditions for the stochastic simulation of a bistable chemical reaction network (purple box) that determine cellular phenotype.

We next derived a multi-scale model that approximates the stochastic growth dynamics of a population endowed with bet-hedging with memory. We explored the selective advantage of memory and whether the evolution of memory and bet-hedging might be predictable. Surprisingly, we found that memory is not universally beneficial and that the sign of the fitness contribution is dependent on both the associated GP-mapping and the genetic configuration of the population. This phenomenon mirrors sign epistasis in population genetics, wherein a mutation can be either beneficial or deleterious dependent on the genetic background in which it occurs. Thus, mirroring our results for fitness landscapes, we find a rugged landscape of bet-hedging mechanisms on which the evolution of bet-hedging need not be repeatable. We infer that knowledge of both genetic and epigenetic inheritance, along with the GP-mapping, is required to make successful evolutionary predictions. To conclude this chapter, we explored the impact of bet-hedging and memory on the efficacy of evolutionary therapies. Critically, we found that despite the difficulty in making evolutionary predictions, therapy can be partially optimised in a mechanism-agnostic way; namely through optimised timing of a metronomic dose that maximises the killing of proliferative cells whilst minimising toxicity.

GP-Mappings at the Population Scale

Throughout our modelling we identified the common theme that successful evolutionary predictions are likely to be predicated on knowledge of the underlying GP-mapping. Unfortunately, where stochasticity drives phenotypic heterogeneity through bet-hedging, this mapping is believed to comprise complex molecular pathways that are difficult to characterise. With this issue in mind, we conclude the thesis by identifying population-scale manifestations of the GP-mappings that can be identified with presently available experimental assays. Specifically, through simulations we identified the phenotype distributions and spectra of mutational frequencies arising from populations endowed with bet-hedging and mutation driven resistance. These models are employed to undertake two case studies: first, we identify memory in bet-hedging as a possible explanation of persister-cells that expand in frequency in confluent populations; and second, we develop a means to detect bet-hedging-driven resistance from matched pre- and post- therapy bulk whole-genome sequencing.

Conclusions

The aim of this thesis was to explore the extent to which evolution can be predicted to optimise evolutionarily-informed therapies. Our work began with the observation that understanding the drivers of phenotypic heterogeneity is key to successful prediction. Through a survey of GP-mapping studies we identified bet-hedging as an under-studied driver of resistance. With this in mind, we first used the fitness landscape model to demonstrate how model GP-mappings can improve evolutionary predictions. Second, we introduced a new model of bet-hedging and demonstrated that the evolutionary dynamics of bet-hedging are closely linked to the specific

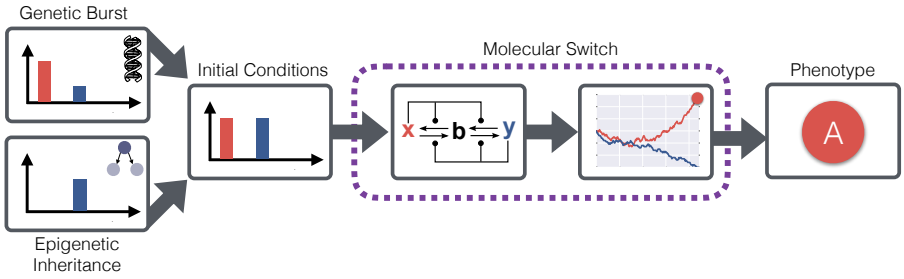

Figure 3: **The extended chemical reaction network model for bet-hedging with memory.** Following division, daughter cells inherit a number of molecules determined by the parental cell age. These molecules are combined with a fixed gene expression profile (genotype) and serve as initial conditions for the bistable chemical reaction network model of the GP–mapping. The inheritance of molecules from the parent cell biases the switch towards the phenotype of the parent, inducing phenotypic memory.

mechanisms through which it arises. Third, we extended our bet-hedging model to account for phenotypic memory, demonstrating that the selective advantage of memory is contingent on the precise machinery of bet-hedging. These findings indicate that the evolution of bet-hedging need not be repeatable, and thus that successful evolutionary predictions will depend on characterisation of both the modes of inheritance and the precise molecular switching. Finally, we undertook two model-driven case studies to demonstrate how the properties of bet-hedging can be identified at the population scale. We anticipate that studies such as these represent the first step towards developing the necessary model systems through which GP-mappings can be studied. Throughout this thesis we have employed abstract models as a means to generate biological hypotheses, but have done so with pragmatic goals; namely, how can evolutionary therapy be improved?

References

[1] Greaves, M. and Maley, C. C. Clonal evolution in cancer. Nature, 481(7381):306–313, 2012.

[2] Gatenby, R. A., Silva, A. S., Gillies, R. J., and Frieden, B. R. Adaptive therapy. Cancer Research, 69(11):4894–4903, 2009.

[3] Pigliucci, M. Genotype–phenotype mapping and the end of the genes as blueprint metaphor. Philosophical Transactions of the Royal Society B: Biological Sciences, 365(1540):557–566, 2010.

[4] Nichol, D., Jeavons, P., Fletcher, A. G., Bonomo, R. A., Maini, P. K., Paul, J. L., Gatenby, R. A., Anderson, A. R., and Scott, J. G. Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. PLoS Computational Biology, 11(9):e1004493, 2015.

[5] Nichol, D., Rutter, J., Bryant, C., Hujer, A. M., Lek, S., Adams, M. D., Jeavons, P., Anderson, A. R., Bonomo, R. A., and Scott, J. G. Antibiotic collateral sensitivity is contingent on the repeatability of evolution. Nature Communications, 10(1):334, 2019.

[6] Lewis, K. Persister cells, dormancy and infectious disease. Nature Reviews Microbiology, 5(1):48–56, 2006.

[7] Cardelli, L. and Csikász-Nagy, A. The cell cycle switch computes approximate majority. Scientific Reports, 2:656, 2012.

[8] Nichol, D., Robertson-Tessi, M., Jeavons, P., and Anderson, A. R. Stochasticity in the genotype-phenotype map: Implications for the robustness and persistence of bet-hedging. Genetics, 204(4):1523–1539, 2016.

[9] Yang, H. W., Chung, M., Kudo, T., and Meyer, T. Competing memories of mitogen and p53 signalling control cell-cycle entry. Nature, 549(7672):404, 2017.

[10] Mira, Portia M and Crona, Kristina and Greene, Devin and Meza, Juan C and Sturmfels, Bernd and Barlow, Miriam Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance PloS one,10(5):e0122283.