A Discrete Nash Theorem with Low Complexity and Dynamic Equilibria

Stéphane Le Roux
LIP, ENS, Lyon, France
JAIST, Nomi, Japan

Pierre Lescanne
LIP, ENS, Lyon, France

René Vestergaard*
JAIST, Nomi, Japan

IS-RR-2006-006
May 31, 2006

*Corresponding author, <vester@jaist.ac.jp>.

A Discrete Nash Theorem with Low Complexity and Dynamic Equilibria

Stéphane Le Roux
LIP, ENS, Lyon, France
JAIST, Nomi, Japan

Pierre Lescanne
LIP, ENS, Lyon, France

René Vestergaard*
JAIST, Nomi, Japan

May 31, 2006

Abstract

Nash's Theorem guarantees the existence of *Nash equilibria* for *strategic-form games*. The typical proof of the result uses Brouwer's Fixed Point Theorem on probabilistic strategies. We show that Tarski's Fixed Point Theorem can be used to establish a similar result for discrete equilibria in a much larger class of games that we call *conversion/preference games*. Our result rests on a graph characterisation of Nash equilibria that i) reifies the decision procedure for *pure* Nash equilibria, ii) allows us to compute the equilibria in quadratic time in the number of game situations, and iii) makes the equilibria explicitly dynamic in nature. We also briefly discuss the extended range of technical applications of *non-cooperative game theory* that results from the new theorem, including for gene regulation and cell-level signal transduction.

1 Introduction

In (1), Nash proved that all *strategic-form games* have a *mixed* (aka probabilistic) Nash equilibrium. A detailed proof using Brouwer's Fixed Point Theorem (2) is given in (3). To be precise, Nash observed that the set of finite strategic-form games can be embedded into a set of (continuous) strategic-form games where each agent's set of strategies is comprised of probability distributions over the agent's original strategies and then proved that the latter set of games always have Nash equilibria, in the form of fixed points of a given continuous function.

	h_1	h_2
v_1	0, 1	1, 0
v_2	1, 0	0, 1

Each cell of the array above should be thought of as the outcome of a possible play of the exemplified game, with the first number in a cell being the resulting *payoff* to player 'v', who chooses the row, and the second number the payoff to player 'h', who chooses the column. The example has no one outcome that satisfies all players in Nash's sense: an agent is happy if he cannot single-handedly improve his payoff, see Definition 2. For example, 'v' would be unhappy with the upper-left outcome because the lower-left is a feasible alternative to him that would increase his payoff; in turn, 'h' would be unhappy with the lower-left outcome because of the lower-right one, and so on counter-clockwise around the array. Instead, a probabilistic Nash equilibrium arises if both agents choose between their two options with equal probability, with *expected* payoffs of a half to each. The usual reading of Nash's probabilistic construction is that it prescribes (weighted) compromises.

In Section 2, we discuss the original Nash Theorem; in Section 3, we introduce the new formalism of conversion/preference (C/P) games; in Section 4,

*Corresponding author, <vester@jaist.ac.jp>.

we pursue a discrete fixed-point construction on C/P games; in Section 5, we directly prove our discrete Nash theorem with low complexity; in Section 6, we show that the new result admits dynamic equilibria; in Section 7, we connect our earlier fixed-point construction to our new Nash theorem; in Section 8, we compare the Nash equilibria found by Nash's probabilistic and our discrete Nash theorems; in Section 9, we briefly discuss the additional modelling capabilities of C/P games compared with Nash's strategic-form games; in Section 10, we summarise two applications of our new Nash theorem in life science.

2 Nash's Theorem

We informally described strategic-form games as being arrays. More formally, we have the following.

Definition 1 (Strategic-Form Games) G^{sf} *are 3-tuples* $\langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, \mathrm{P}\rangle$*, where:*

- $\mathcal{A}$ *is a non-empty set of* agents,
- $\mathrm{S}_{\mathcal{A}}$ *is the cartesian product,* $\bigotimes_{\mathrm{a}\in\mathcal{A}} \mathrm{S}_{\mathrm{a}}$*, called* strategy profiles, *of non-empty sets of* individual strategies, S_{a}*, for each agent,* a,
- $\mathrm{P} : \mathrm{S}_{\mathcal{A}} \to \mathcal{A} \to \mathbb{R}$*, is a real-valued* payoff function.

Let s*, a strategy profile, range over* $\mathrm{S}_{\mathcal{A}}$ *and let* s_{a} *be the* a*-projection of* s.

In terms of Nash equilibria, agents in a strategic-form game are free to change the entry in their dimension of the cartesian product of individual strategies but must leave any other entries unchanged. The question of whether a better outcome single-handedly can be obtained by any particular agent is therefore answered (in the negative) by the following predicate.

Definition 2 *A strategy,* s*, is a* Nash equilibrium, $\mathrm{Eq}^{\mathrm{N}}_{\mathrm{G}^{\mathrm{sf}}}(\mathrm{s})$*, for a given strategic-form game,* G^{sf}*, iff*[1]

$$\begin{array}{lll} \forall \mathrm{a} \in \mathcal{A}, \mathrm{s}' \in \mathrm{S}_{\mathcal{A}} & . & (\forall \mathrm{a}' \in \mathcal{A} \,.\, \mathrm{a} \neq \mathrm{a}' \Rightarrow \mathrm{s}_{\mathrm{a}'} = \mathrm{s}'_{\mathrm{a}'}) \\ & & \Downarrow \\ & & \neg(\mathrm{P}(\mathrm{s})(\mathrm{a}) < \mathrm{P}(\mathrm{s}')(\mathrm{a})) \end{array}$$

[1]Standard notation for our s′ is s; σ_{a}, i.e., s with something else in the a-position.

We suppress the G^{sf}-subscript when it is obvious from the context what strategic-form game we mean.

As noted, Nash equilibria do not always exist directly in a strategic-form game and we now define *individual probabilities*, *probability profiles*, and, for a given probability profile, the overall *probability* that the agents collectively assign to a strategy profile.

Definition 3 (Strategic-Form Probabilities) *Given finite* $\langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, \mathrm{P}\rangle$.

$$\begin{aligned} \mathrm{W}^{\mathrm{S}_{\mathrm{a}}} &\triangleq \{\mathrm{w}_{\mathrm{a}} : \mathrm{S}_{\mathrm{a}} \to [0,1] \mid (\textstyle\sum_{\sigma\in \mathrm{S}_{\mathrm{a}}} \mathrm{w}_{\mathrm{a}}(\sigma)) = 1\} \\ \mathrm{W}^{\mathrm{S}_{\mathcal{A}}} &\triangleq \textstyle\bigotimes_{\mathrm{a}\in\mathcal{A}} \mathrm{W}^{\mathrm{S}_{\mathrm{a}}} \\ \mu^{\mathrm{S}_{\mathcal{A}}}(\mathrm{w},\mathrm{s}) &\triangleq \textstyle\prod_{\mathrm{a}\in\mathcal{A}} \mathrm{w}_{\mathrm{a}}(\mathrm{s}_{\mathrm{a}}) \end{aligned}$$

The *expected payoff function* associated with a probability profile is as follows.

Definition 4 (Expected Payoff Function) *Given finite* $\langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, \mathrm{P}\rangle$ *with associated probabilities.*

$$\mathrm{EP}^{\mathrm{S}_{\mathcal{A}}}_{\mathrm{P}}(\mathrm{w})(\mathrm{a}) \quad \triangleq \quad \textstyle\sum_{\mathrm{s}\in \mathrm{S}_{\mathcal{A}}} \mu^{\mathrm{S}_{\mathcal{A}}}(\mathrm{w},\mathrm{s}) \cdot \mathrm{P}(\mathrm{s})(\mathrm{a})$$

Nash's result, existence argument, and the employed construction can now be articulated as follows.

Theorem 5 (Nash (1, 3)) *Consider a finite strategic-form game,* $\langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, \mathrm{P}\rangle$*, given with probabilities. The strategic-form game* $\langle \mathcal{A}, \mathrm{W}^{\mathrm{S}_{\mathcal{A}}}, \mathrm{EP}^{\mathrm{S}_{\mathcal{A}}}_{\mathrm{P}}\rangle$ *has a Nash equilibrium. Informally, we say that* $\langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, \mathrm{P}\rangle$ *has a* probabilistic Nash equilibrium.

Proof We follow (3). $\mathrm{W}^{\mathrm{S}_{\mathcal{A}}}$ is the cartesian product of each agent's $\mathrm{W}^{\mathrm{S}_{\mathrm{a}}}$. Because they involve a summation to 1, each $\mathrm{W}^{\mathrm{S}_{\mathrm{a}}}$ is the standard simplex of the vector space spanned by the elements of S_{a} taken as unit vectors. As a result, $\mathrm{W}^{\mathrm{S}_{\mathcal{A}}}$ is a convex polytope in the vector space spanned by $\mathrm{S}_{\mathcal{A}}$, which in particular means that it is compact. More, it is possible to define a continuous function from probability profiles to probability profiles that, for each agent, speculatively puts more weight where that agent can benefit from it *relative to the other agents' unchanged weights* and then re-normalises all weights and makes a combined change. This function has a fixed point by the generalised Brouwer's Fixed Point Theorem[2] (2) and any such fixed point is a Nash equilibrium (3). □

[2]"A continuous function on a compact, convex set has a fixed point".

The problem of finding probabilistic Nash equilibria for finite strategic-form games with at least two agents is in PPAD in the size of $\mathrm{S}_{\mathcal{A}}$ (4). In fact, it is PPAD-complete (5, 6). Informally, this means that it is likely that we will not know whether the problem is polynomial or exponential for some time.

3 C/P Games

To facilitate our discrete development, we now introduce a new game formalism called conversion/preference (C/P) games. It is based on strategic-form games, with conversion and preference accounting for the views and options available to the partaking players. The formalism distinguishes itself by seemingly being the most general structure that allows for the definition of Nash equilibria; in a sense, C/P games capture the essence of Nash equilibria.

Definition 6 (C/P Games) G^{cp} *are 4-tuples* $\langle \mathcal{A}, \mathcal{S}, (\underset{\mathrm{a}}{\succ\!\!-})_{\mathrm{a}\in\mathcal{A}}, (\lhd_{\mathrm{a}})_{\mathrm{a}\in\mathcal{A}}\rangle$, *where:*

- $\mathcal{A}$ *is a non-empty set of* agents.
- $\mathcal{S}$ *is a non-empty set of* synopses.[3]
- *For* $\mathrm{a} \in \mathcal{A}$, $\underset{\mathrm{a}}{\succ\!\!-}$ *is a binary relation over* $\mathcal{S}$, *associating two synopses if agent* a *can* convert *the first to the second.*
- *For* $\mathrm{a} \in \mathcal{A}$, $\lhd_{\mathrm{a}}$ *is a binary relation over* $\mathcal{S}$, *associating two synopsis if agent* a prefers *the second to the first.*

Synopses are abstractions over *strategy profiles* but, in combination with conversion and preference, they can also be seen more generally as denoting the state of (a play of) a game, e.g., in terms of the players' possessions or even in a purely abstract sense, see Sections 9 and 10.

In strategic-form games, the cartesian-product nature of the set of strategy profiles, $\mathrm{S}_{\mathcal{A}}$, determines what alternatives are available to a particular agent as far as Nash equilibria are concerned: agent a can change the a-projection of an s to something else from S_{a}. The dynamic view of an agent "changing" one strategy profile to another by contributing a different individual strategy is what we capture directly in the definition of C/P games, by $\underset{\mathrm{a}}{\succ\!\!-}$, without resorting to an underlying structure of, in this case, $\mathcal{S}$.

The preference relations, $\lhd_{\mathrm{a}}$, account relatively for the penalties and rewards the agents receive in the different outcomes, without explicit payoffs (7).

Definition 7 *A synopsis,* s, *is an* abstract Nash equilibrium, $\mathrm{Eq}^{\mathrm{aN}}_{\mathrm{G}^{\mathrm{cp}}}(\mathrm{s})$, *for a given* G^{cp}, *if*

$$\forall \mathrm{a} \in \mathcal{A}, \mathrm{s}' \in \mathcal{S} \quad . \quad \mathrm{s} \underset{\mathrm{a}}{\succ\!\!-} \mathrm{s}' \;\Rightarrow\; \neg(\mathrm{s} \lhd_{\mathrm{a}} \mathrm{s}')$$

We suppress the G^{cp}-subscript when it is obvious from the context what C/P game we mean.

Unsurprisingly, the canonical embedding of strategic-form games into C/P games preserves and reflects (abstract) Nash equilibria and, consequently, we typically suppress the word "abstract".

Theorem 8 *For a strategic-form game,* G^{sf}, *let*

$$\begin{array}{rcl} \mathrm{s} \underset{\mathrm{a}}{\succ\!\!-\!\!\prec} \mathrm{s}' & \triangleq & \forall \mathrm{a}' \,.\, (\mathrm{a} \neq \mathrm{a}' \Rightarrow \mathrm{s}_{\mathrm{a}'} = \mathrm{s}'_{\mathrm{a}'}) \\ \mathrm{s} \lhd_{\mathrm{a}} \mathrm{s}' & \triangleq & \mathrm{P}(\mathrm{s})(\mathrm{a}) < \mathrm{P}(\mathrm{s}')(\mathrm{a}) \end{array}$$

$\mathrm{G}^{\mathrm{cp}} = \langle \mathcal{A}, \mathrm{S}_{\mathcal{A}}, (\underset{\mathrm{a}}{\succ\!\!-\!\!\prec})_{\mathrm{a}\in\mathcal{A}}, (\lhd_{\mathrm{a}})_{\mathrm{a}\in\mathcal{A}}\rangle$ *is a C/P game and*

$$\mathrm{Eq}^{\mathrm{N}}_{\mathrm{G}^{\mathrm{sf}}}(\mathrm{s}) \;\Leftrightarrow\; \mathrm{Eq}^{\mathrm{aN}}_{\mathrm{G}^{\mathrm{cp}}}(\mathrm{s})$$

To illustrate, we note that the example we considered earlier is mapped to the following C/P game.

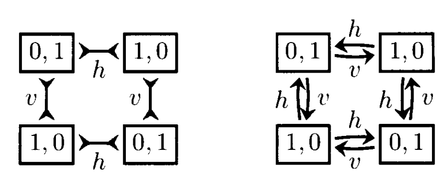

The C/P-game formalism accommodates more than just strategic-form games. We shall address what this means in Section 9 and take more substantial advantage of it in Section 10.

[3]The name synopsis is inspired by the thespian meaning of 'abstract of a play'.

$$\frac{s \underset{a}{\succ} s' \quad s \lhd_a s'}{s \rightarrow_a s'}$$

Fig. 1. The (free) a-*change-of-mind* relation

$$\frac{s \rightarrow_a s' \quad \lfloor s \rfloor \neq \lfloor s' \rfloor}{\lfloor s \rfloor \curvearrowright_a \lfloor s' \rfloor}$$

Fig. 2. The a-*progressive change-of-mind* relation

4 Discrete Fixed Points

The starting point of our discrete Nash Theorem is the relation that captures when agents both *can* and *want* to change their minds.

Definition 9 *Given* G^{cp}, *the* (free) change-of-mind *relation,* $\rightarrow_a$, *for agent* a *is given in Figure 1. Let* $\rightarrow \triangleq \bigcup_{a \in \mathcal{A}} \rightarrow_a$ *and let* $\rightarrow^*$ *be the reflexive, transitive closure of* $\rightarrow$.

Our initial interest in the change-of-mind relation is that it allows us to define a discrete equivalent of the probability-update function sketched in the proof of Theorem 5. Rather than trying to establish that this function is of inherent/instinctive interest as, for example, Nash's is thought to be, we shall objectively compare the notions of compromise induced by the two different flavours of Nash equilibria in Section 8.[4]

Definition 10 (Discrete Updates)

$$\mathscr{U}(S) \triangleq \bigcup_{s \in \mathcal{S}} \{s' \in S \mid s \rightarrow^* s'\}$$

The empty set, $\emptyset$, and the full set, $\mathcal{S}$, are straightforwardly seen to be fixed points of this function. More generally, we have the following result.

Lemma 11 *The set of fixed points of* $\mathscr{U}$ *forms a non-empty, complete lattice.*

Proof $\mathscr{U}$ is monotonic on the complete lattice $\mathcal{P}(\mathcal{S})$ ordered by inclusion because $\rightarrow^*$ is reflexive; we are done by Tarski's Fixed Point Theorem (8). □

[4]Additionally, Section 10.1 will detail a compelling analogue of the difference between the two styles of update functions that is well-established in the area of gene-regulation analysis.

Unlike Nash, our fixed-point construction works natively for infinite games and, of course, for games with unstructured/arbitrary $\mathcal{S}$. On the other hand, not all fixed points will be Nash equilibria and we therefore pursue an alternative presentation of our technology next. We shall return to the fixed-point construction and the Nash-equilibrium characterising equivalence proof of the two approaches in Section 7.

5 A Discrete Nash Theorem

In this section, we take inspiration from the fact that the change-of-mind relation can be used to characterise Nash equilibria in the pure sense. We note that a terminal in a graph is a node with no out-edges.

Proposition 12 $Eq^{aN}(s) \Leftrightarrow Terminal_{\rightarrow}(s)$

Proof Straightforward, as $\rightarrow_a$ is the intersection of $\underset{a}{\succ}$ and $\lhd_a$. □

The result says that the notion of "happiness" discussed in Section 1 corresponds to the absence of change-of-mind steps and it implies that only cycles can prevent the existence of Nash equilibria for a finite C/P game. Thinking graph theoretically, we note that the general form of cycles in graphs is called *strongly connected components* (SCCs) and that their cycle-free super-structure is called the *shrunken graph*, see Appendix A.

Definition 13 (Progressive Change-of-Mind) *The* progressive change-of-mind *relation,* $\curvearrowright_a$, *for agent* a *is given in Figure 2. Let* $\curvearrowright \triangleq \bigcup_{a \in \mathcal{A}} \curvearrowright_a$.

The progressive change-of-mind relation is the shrunken graph of the change-of-mind relation, which means that it is cycle-free by construction. It also means that we are able to prove a first version of our alternative Nash Theorem, establishing the guaranteed existence of a discrete notion of Nash equilibria.

Theorem 14 *Consider a finite C/P game. The "shrunken" C/P game* $\langle \mathcal{A}, \lfloor \mathcal{S} \rfloor, (\curvearrowright_a)_{a \in \mathcal{A}}, (\curvearrowright_a)_{a \in \mathcal{A}} \rangle$ *has Nash equilibria and all of them can be found in quadratic time in the size of* $\mathcal{S}$.

Proof The progressive change-of-mind relation underlying the considered "shrunken" C/P game is cycle-free by construction, see Proposition 20, Appendix A. This means that the length of any change-of-mind path in it is bounded by the size of $\lfloor \mathcal{S} \rfloor$, guaranteeing the existence of terminal nodes and thus, by Proposition 12, compromises that are "shrunken" Nash equilibria. The complexity measure is due to Tarjan (9), see Theorem 19, Appendix A. □

The theorem exploits what could be called the *Nash Construction*, i.e., the guarantee first used in Theorem 5 that some derived game has Nash equilibria that are meaningful as constituting a compromise when retracted back to the original game. We will spend the latter part of this article, from Section 8 on-wards, clarifying the notion of compromise that is invoked in Theorem 14 and on comparing it to Nash's payoff-driven probabilistic notion.

6 Change-of-Mind Equilibria

Unlike Nash's Theorem, the compromises fingered as Nash equilibria in Theorem 14 have a formal, direct characterisation in the originating C/P game. As we shall see, they are clusters of synopses that, while potentially improvable in the view of some agents, cannot be improved upon in an irreversible manner.

Definition 15 *Write* $\xrightarrow{\mathrm{S}}$ *for* $\rightarrow \cap (\mathrm{S} \times \mathrm{S})$. *The graph of an SCC of synopses,* $\xrightarrow{\lfloor s \rfloor}$, *is a* change-of-mind equilibrium, $\mathrm{Eq}^{\mathrm{com}}_{\mathrm{G^{cp}}}(\xrightarrow{\lfloor s \rfloor})$, *for a given C/P game,* $\mathrm{G^{cp}}$, *if*

$$\forall s_0 \in \lfloor s \rfloor, s' \in \mathcal{S} \quad . \quad s_0 \rightarrow^* s' \ \Rightarrow \ s' \in \lfloor s \rfloor$$

As before, we will suppress the $\mathrm{G^{cp}}$-subscript when it is obvious what C/P game we mean. The point is this: a change-of-mind equilibrium is the topology in the original game of the discrete compromises fingered as Nash equilibria in Theorem 14.

Lemma 16 $\mathrm{Eq}^{\mathrm{com}}_{\mathrm{G^{cp}}}(\xrightarrow{\lfloor s \rfloor}) \quad \Leftrightarrow \quad \mathrm{Eq}^{\mathrm{aN}}_{\lfloor \mathrm{G^{cp}} \rfloor}(\lfloor s \rfloor)$

Proof By two direct arguments. □

Unlike Theorems 5 and 14, whose formulations depend on a derived game, we can now state our main theorem, establishing the existence of a dynamic notion of Nash equilibria directly in the game concerned.

Theorem 17 *Any finite C/P game,* $\mathrm{G^{cp}}$, *has change-of-mind equilibria,* $\mathrm{Eq}^{\mathrm{com}}_{\mathrm{G^{cp}}}$, *and all of them can be found in quadratic time in the number of synopses.*

Proof Lemma 16 and Theorem 14. □

To preliminarily illustrate what this result means, we revisit our motivating example from Section 1.

	h_1	h_2
v_1	0, 1	1, 0
v_2	1, 0	0, 1

$$\begin{array}{ccc} 0,1 & \leftarrow & 1,0 \\ \downarrow & & \uparrow \\ 1,0 & \rightarrow & 0,1 \end{array}$$

As mentioned, the only probabilistic Nash equilibrium arises when both agents choose between their two options with equal probability, for expected payoffs of a half to each. The only change-of-mind equilibrium is shown on the right. It, too, involves all four outcomes of the game in the prescribed compromise. The main virtue of Nash's probabilistic compromise is that it dictates an exact expected payoff to each agent. Change-of-mind equilibria, on the other hand, states exactly why the four outcomes are included in the compromise. The upper-left outcome, say, is included because 'h' prefers it to the upper-right outcome, and so on clock-wise around the array.

7 Fixed Points Revisited

Returning to our discrete fixed-point construction in Section 4, we can now finally identify the interesting $\mathscr{U}$-fixed points as those that are least non-empty.

Lemma 18 *Given* $\mathrm{G^{cp}}$ *with change-of-mind* $\rightarrow$.

$$\mathrm{Eq}^{\mathrm{com}}(\xrightarrow{\lfloor s \rfloor})$$
$$\Updownarrow$$
$$\mathscr{U}(\lfloor s \rfloor) = \lfloor s \rfloor \ \wedge \ (\forall \mathrm{S} . \emptyset \subsetneq \mathrm{S} \subsetneq \lfloor s \rfloor \Rightarrow \mathscr{U}(\mathrm{S}) \nsubseteq \mathrm{S})$$

Proof The "⇑" part follows by a direct argument, while the other direction is split into two lemmas, each proved by a direct argument.

- $\mathrm{Eq}^{\mathrm{com}}(\xrightarrow{\lfloor s \rfloor}) \Rightarrow \mathscr{U}(\lfloor s \rfloor) = \lfloor s \rfloor$
- $\emptyset \subsetneq \mathrm{S} \subsetneq \lfloor s \rfloor \Rightarrow \mathscr{U}(\mathrm{S}) \not\subseteq \mathrm{S}$ □

The result is, perhaps, most interesting in the infinite case where we have not guaranteed the existence of change-of-mind equilibria. Some infinite C/P games may have change-of-mind equilibria but, e.g., any C/P game with the following infinitely-ascending progressive change-of-mind relation will not.

Clearly, the relation has no terminal node. Conversely, the complete lattice of $\mathscr{U}$-fixed points is the set of all tails in which there is no least non-empty element. The relevance of this, however, is that it may open a way to define a notion of Nash equilibrium that specialises to that of Definition 7 in the finite case and can be guaranteed to exist also in the infinite case, in case that should be needed.

8 Compromises

Our running example strategic-form game betrays the substantial differences that may exist between the compromises prescribed by Nash's probabilistic Theorem 5 and our discrete Theorems 14 and 17.

8.1 Direct Comparisons

In the example, the compromises coincide, i.e., the only probabilistic Nash equilibrium assigns non-zero probabilities to the same outcomes as are involved in the only change-of-mind equilibrium. We will now show that all possible configurations can be observed when comparing compromises.

Change-of-Mind are Probabilistic Generalising our running example to a three-by-three game highlights perhaps the most interesting feature of change-of-mind equilibria, namely the ability to carve-out a part of a game as constituting a Nash equilibrium, below right with six involved outcomes.

	h_1	h_2	h_3
v_1	0, 1	0, 0	1, 0
v_2	1, 0	0, 1	0, 0
v_3	0, 0	1, 0	0, 1

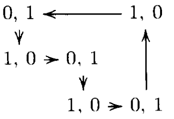

The only probabilistic Nash equilibrium arises again if both agents choose between their options with equal probability, for expected payoffs of a third to each and involving all nine possible outcomes.

Probabilistic are Change-of-Mind A different generalisation of our two-by-two example arises by adding an extra, 'h'-undesirable column.

	h_1	h_2	h_3
v_1	0, 1	1, 0	0, 1
v_2	1, 0	0, 1	1, −7

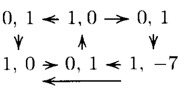

If 'h' puts weight on the third column, 'v' will prefer the second row because of the double 1 payoff for him there. This, however, will give 'h' a negative payoff and 'h' therefore wishes to avoid h_3, i.e., the only probabilistic Nash equilibrium involves v_1, v_2, h_1, and h_2 with equal probabilities. The one change-of-mind equilibrium involves all six outcomes.

Disjoint Compromises By a similar token, we can make several rows 'v'-undesirable.

	h_1	h_2	h_3
v_1	0, 1	−7, 0	1, 0
v_2	1, 0	0, 1	−7, 0
v_3	−7, 0	1, 0	0, 1
v_4	0, 0	0, 0	0, 0

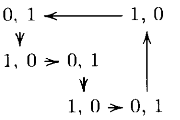

In any Nash probabilistic compromise, player 'v' chooses strategy v_4 with full weight. By contrast, the only change-of-mind equilibrium is disjoint from there, involving the previously-observed cycle around the cells with 1, 0 and 0, 1.

Non-Trivial Overlaps The strategic-form game where only the last row is 'v'-undesirable exhibits complementary features.

	h_1	h_2	h_3
v_1	0, 1	0,0	1,0
v_2	1, 0	0,1	0,0
v_3	0, 0	1,0	0,1
v_4	−7,0	1,0	0,1

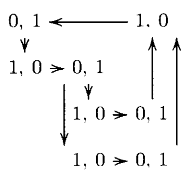

As before, 'v' will avoid the row with the negative payoff and the only probabilistic Nash equilibrium involves the upper nine cells with equal probability. The only change-of-mind equilibrium is as shown.

8.2 Absence of Objective Measures

Focusing narrowly on their prescribed compromises for strategic-form games, we have not been able to separate change-of-mind and probabilistic Nash equilibria in terms of a measure. In particular, the examples in Section 8.1 make it clear that neither notion consistently results in smaller compromises than the other nor higher average/expected payoffs. Indeed, the two notions appear to essentially be of independent interest. More, the algebraic argument we referred to as Nash's construction is sufficiently general to allow for other kinds of derived Nash equilibria.

9 C/P vs Strategic-Form

The prisoner's dilemma is a classic example in game theory that more than being a dilemma highlights the non-cooperative aspect of Nash equilibria. The game is made up of two agents who are accused of a crime. If both confess, they share the mandated 3-year prison term for the crime, with right of parole. If one confesses and the other denies any involvement, the former serves the full sentence. If both deny their involvement (and they have been found guilty), they are deemed to have attempted to pervert the course of justice and are punished for that and the crime.

	h_{confess}	h_{deny}
v_{confess}	−1, −1	−3, 0
v_{deny}	0, −3	−2, −2

The issue at hand is that the lower right cell, with both denying their involvement, is the only Nash equilibrium although they would be better off by jointly confessing. The mechanics of the prisoner's dilemma also show up if we consider two players that share two tokens that they play for: a player with a token can take the token from the other player at will with the aim of acquiring both tokens. We call the game "blink-and-you-lose".

	h_{leave}	h_{take}
v_{leave}	1, 1	0, 2
v_{take}	2, 0	1, 1

Here, the lower-right cell is again the only Nash equilibrium. The main failure of non-cooperative game theory in this case is not that it steers the players towards the worse of two evils but that it fails to allow for the winning configurations of one player possessing both tokens (corresponding to avoiding incarceration in the prisoner's dilemma) as equilibria.

By contrast, and because C/P games need not be array-shaped and because no particular set of permissible moves are mandated, we can model aspects of blink-and-you-lose that classic game theory cannot. One way is to form a game consisting not of four strategy profiles but of three "game situations": a) player 1 has both tokens, b) the players have a token each, and c) player 2 has both tokens. According to the rules of the game, agent 1 prefers his winning situation, a), to b) and c), and the neutral b) to c), his losing situation. Similarly, for agent 2 and c) over a) and b), and b) over a). When it comes to specifying the conversion relations there are at least three distinct principles that can be employed, i.e., agents can be assigned different capabilities/intentions.

Foresight: A player realises that he can win by taking the opponent's token faster than the opponent can react, i.e., player 1 can convert b) to

a) by outpacing player 2. Player 2, in turn, can convert b) to c). This version of the game has two singleton change-of-mind equilibria, i.e., Nash equilibria: a) and c).

$$(2,0) \xleftarrow[1]{} (1,1) \xrightarrow[2]{} (0,2)$$

Hindsight: A player, say 1, analyses what would happen if he does not act. In case 2 acts, the game would end up in c) and 1 looses, and 1 therefore concludes that he could have prevented the c) outcome by acting. In other words, it is within 1's power to convert c) to b). Similarly for player 2 from a) to b). This version of the game has one singleton change-of-mind equilibrium, i.e., a Nash equilibrium: b).

$$(2,0) \xrightarrow[2]{} (1,1) \xleftarrow[1]{} (0,2)$$

Omnisight: The players have both conversion step-styles just considered, resulting in a C/P game with one change-of-mind equilibrium covering all outcomes, i.e., no *pure* Nash equilibrium.

$$\begin{tikzcd} (2,0) \arrow[r, bend right, "2"'] & (1,1) \arrow[l, bend right, "1"'] \arrow[r, bend left, "2"] & (0,2) \arrow[l, bend left, "1"] \end{tikzcd}$$

10 Rewriting Game Theory

As mentioned, our main result is Theorem 17. We have dubbed applications of C/P games and change-of-mind equilibria *rewriting game theory* to stress our positive reading of the change-of-mind relation that Nash interpreted negatively, see Proposition 12. Graphs are the simplest and most abstract structure pursued in the field of rewriting, with focus on their dynamic aspects, equational theories, termination properties, and more (10).

We now present the highlights of two uses of rewriting game theory. Generally speaking, they show that change-of-mind equilibria can be used to identify the biologically meaningful parts of a purely chemically-conceived C/P game.

10.1 Gene Regulation

In (11), we use rewriting game theory to provide a foundation for Kauffman/Thomas-style gene-regulation analysis (12–15). In the case of the standard example of 2-variable *bacteriophage lambda*, Kauffman's model produces the following state-space graph, which we directly take as change-of-mind. The graph consists of *synchronous* updates of gene states.

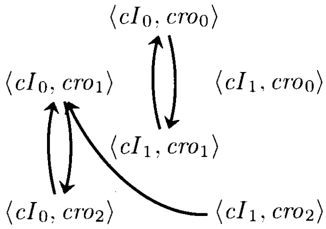

Thomas' model, on the other hand, is *asynchronous*, allowing each gene to update its own state in isolation, with the dotted arrows indicating updates of *cro* and the dashed arrows updates of *cI*.[5]

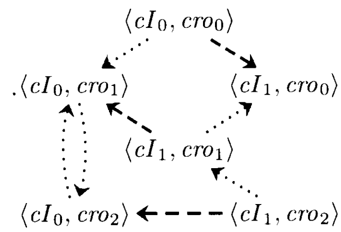

The relevance of Kaufmann/Thomas-style gene-regulation analysis is that, e.g., the singleton change-of-mind equilibrium in these graphs consisting of $\langle cI_1, cro_0\rangle$ is phage λ's *lysogenic* state that "involves integration of the phage DNA into the bacterial chromosome [of its host] where it is passively replicated at each cell division — just as though it were a legitimate part of the bacterial genome" (16). Similarly, the change-of-mind equilibrium consisting of the cycle between $\langle cI_0, cro_1\rangle$ and $\langle cI_0, cro_2\rangle$ in

[5] By analogy to our's vs Nash's game-theoretic update functions, we see that Nash and Kauffman correspond to each other, while we natively correspond to the Thomas approach; we can, however, vary the considered set of agents without otherwise changing a C/P game, which means that we can accommodate also the former two, as already illustrated.

the graphs is characteristic of phage λ's *lytic* state in which it actively uses its host's transcription mechanism to replicate itself (16).[6]

Rewriting game theory provides a unifying account of *static* and *dynamic steady states* of genes as change-of-mind equilibria and, moreover, uniformly accommodates both Kauffman's and Thomas' models and, for that matter, any hybrid of them. In addition, we prove equivalent two independent characterisations of dynamic steady states (as fixed-points (17) and as terminal SCCs (18)), see Lemma 18. Rewriting game theory is a good foundation for gene-regulation analysis because it provides a general theory of possibly dynamic equilibria and because the technical means by which it accomplishes this, i.e., change-of-mind, has the same reading as the established models but for gene-independent reasons.

10.2 Signalling Pathways

In (19), we construct a so-called *cascaded protein game* from the 113 chemical reactions stated to be involved in MAPK cascades in (20–27). The constructed game has two change-of-mind equilibria. One is the ERK pathway, known as the signalling pathway responsible for cell growth, while the other is a combination of the JNK and p38 signalling pathways, known among other things for their cross-talk. The (C/P) agents in the constructed cascaded protein game are the enzymes that catalyse the involved reactions (for a standard increase in observed kinetics of 10^6 to 10^{12} times (28)). The reading of this is i) that the discussed signalling pathways are *inevitable*, i.e., they are the best compromise for what the enzymes prefer to do when given a suitable input, see Theorem 14, and ii) that the pathways are good candidates for a central building block of a living organism because they are *sustainable*, i.e., no enzyme can defect "play" from the pathways once it has arrived there, see Theorem 17.

[6]The cycle between $\langle cI_0, cro_0\rangle$ and $\langle cI_1, cro_1\rangle$ is a known false positive of Kauffman's model.

11 Conclusion

We have proved a new Nash Theorem in two versions: Theorem 14 is stated in the usual indirect form with Nash equilibria in a derived game while Theorem 17 is direct and pertains to the new notion of change-of-mind equilibria. Summarising our development compared with Nash's, we have arbitrary vs array-structured games, arbitrary vs real-valued payoffs, Tarski's vs Brouwer's Fixed Point Theorems, quadratic vs PPAD-complete complexity, and discrete vs probabilistic and dynamic vs payoff-driven equilibria. The result allows for an extended range of technical applications of non-cooperative game theory, e.g., in the life sciences, see Section 10.

A SCCs, Shrunken Graphs

- A *graph* is a binary relation on a carrier set, called vertices: $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$.
- The *reflexive, transitive (or pre-order) closure*, $\rightarrow^*$, of a graph, $\rightarrow$, is

$$\frac{v_1 \rightarrow v_2}{v_1 \rightarrow^* v_2} \qquad \frac{}{v \rightarrow^* v} \qquad \frac{v_1 \rightarrow^* v \quad v \rightarrow^* v_2}{v_1 \rightarrow^* v_2}$$

- The strongly connected component (SCC) of a vertex, v, in a graph is

$$\lfloor v \rfloor \quad \triangleq \quad \{v' \mid v \rightarrow^* v' \wedge v' \rightarrow^* v\}$$

  (The relation "is in the $\lfloor - \rfloor$-class of" is an equivalence relation.)
- The set of SCCs of a graph is

$$\lfloor \mathcal{V} \rfloor \quad \triangleq \quad \{\lfloor v \rfloor \mid v \in \mathcal{V}\}$$

- The *shrunken graph* of $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$ is $\curvearrowright \subseteq \lfloor \mathcal{V} \rfloor \times \lfloor \mathcal{V} \rfloor$, defined by

$$V_a \curvearrowright V_b \triangleq V_a \neq V_b \wedge (\exists v_a \in V_a, v_b \in V_b \,.\, v_a \rightarrow v_b)$$

Theorem 19 (Tarjan (9, 29)) *Given a graph, $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$, the SCCs and their shrunken graph can be found in linear time in the sizes of $\rightarrow$ and $\mathcal{V}$.*

Proposition 20 *A shrunken graph is cycle-free.*

References

1. Nash, John. (1950) *Proceedings of the National Academy of Sciences* **36**, 48–49.
2. Brouwer, L. E. J. (1912) *Mathematische Annalen* **71**, 97–115.
3. Nash, John. (1951) *Annals of Mathematics* **54**.
4. Papadimitriou, C. H. (1994) *Journal of Computer and System Sciences* **48**, 498–532.
5. Chen, Xi & Deng, Xiaotie. (2005) *Elec. Coll. on Comp. Compl. (140)*.
6. Daskalakis, K., Goldberg, P. W. & Papadimitriou, C. H. (2005) *Elec. Coll. on Comp. Compl. (115)*.
7. Osborne, M. J. & Rubinstein, A. (1994) *A Course in Game Theory.* (The MIT Press, Cambridge, Massachusetts).
8. Tarski, Alfred. (1955) *Pacific Journal of Mathematics* **5**, 285–309.
9. Tarjan, Robert E. (1972) *SIAM Journal on computing* pp. 146–160.
10. Terese. (2003) *Term Rewriting Systems*, Cambridge Tracts in Theoretical Computer Science. (Cambridge University Press).
11. Chettaoui, C., Delaplace, F., Lescanne, P., Vestergaard, M. & Vestergaard, R. (2006) Rewriting game theory as a foundation for state-based models of gene regulation, (JAIST), Research Report IS-RR-2006-008.
12. Kauffman, Stuart A. (1969) *Journal of Theoretical Biology* **22**, 437–467.
13. Kauffman, Stuart A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution.* (Oxford University Press).
14. Thomas, René. (1973) *Journal of Theoretical Biology* **42**, 563–585.
15. Thomas, René, Thieffry, Denis & Kaufman, Marcelle. (1995) *Bulletin of Mathematical Biology* **57**, 247–276.
16. Watson, James D., Baker, Tania A., Bell, Stephen P., Gann, Alexander, Levine, Michael & Losick, Richard. (2004) *Molecular Biology of the Gene, 5th edition.* (The Benjamin/Cummings Publishing Company).
17. Thomas, R. & Kaufman, M. (2001) *Chaos* **11**, 180–195.
18. Chaouiya, Claudine, Remy, Elizabeth & Thieffry, Denis. (2005) *Proceedings of CMSB-3.*
19. Senachak, J., Vestergaard, M. & Vestergaard, R. (2006) Rewriting game theory applied to protein signalling in MAPK cascades, (JAIST), Research Report IS-RR-2006-007.
20. BioCarta. (May, 2006) BioCarta. `http://www.biocarta.com`.
21. Chen, Z. et al. (2001) *Chemical Review* **101**, 2449–2476.
22. Kanehisa, M. & associates. (May, 2006) KEGG, MAP signaling pathway - homo sapiens (human). `http://www.genome.jp/kegg/`.
23. Kiriakis, J.M. & Avruch, J. (2001) *Physiological Reviews* **81**.
24. Oda, K., Matsuoka, Y., Funahashi, A. & Kitano, H. (2005) *Molecular Systems Biology.*
25. ProteinLounge. (May, 2006) Proteinlounge databases. `http://www.proteinlounge.com`.
26. Qi, M. & Elion, E.A. (2003) *Journal of Cell Science* **116**, 3569–3572.
27. Upstate. (May, 2006) Cell signaling, MAPK signaling pathway. `http://www.cellsignaling.com`.
28. Voet, Donald & Voet, Judith G. (1995) *Biochemistry.* (John Wiley and Sons, Inc).
29. Mehlhorn, & Naher, . (1995) *CACM: Communications of the ACM* **38**.