

A discrete Nash Theorem with quadratic complexity and dynamic equilibria

Stéphane Le Roux
ENS de Lyon

Pierre Lescanne
ENS de Lyon

René Vestergaard*
JAIST

July 13, 2007

Abstract

Nash's Theorem guarantees the existence of compromising *Nash equilibria* for finite *strategic games*. The typical proof of the result uses Brouwer's Fixed Point Theorem on probabilistic strategies. We show that Tarski's Fixed Point Theorem points to a similar result, except with discrete equilibria and for a larger class of games that we call *conversion/preference games*: C/P games, for short. Our result rests on a unifying graph characterisation of our and Nash's equilibria that, additionally, i) reifies the obvious decision procedure for pure Nash equilibria, ii) allows us to compute our compromising equilibria in quadratic time in the number of game situations, and iii) makes our compromising equilibria explicitly dynamic in nature. We conclude by discussing key examples and the extended range of applications of *Nash-style game theory* that our result enables.

Keywords: Nash's Theorem, equilibria, compromise, behaviour, graphs.

1 Introduction

We revisit Nash's Theorem guaranteeing the existence of certain *Nash equilibria* [12]. As it turns out, Nash's development involves choices that can be

*Corresponding author: René Vestergaard, JAIST, 1-1 Asahidai, Nomi, Ishikawa 923-1292, Japan; <vester@jaist.ac.jp>.

made differently. The main choice we make differently is to pursue a *powerset-game* construction, rather than the probabilistic construction that is normally used to underpin the necessary notion of compromise between strategies. Technically speaking, this means that we will be employing Tarski’s rather than Brouwer’s Fixed Point Theorem in our *update-function* characterisation of the equilibria. Interestingly, this choice makes our fixed-point argument cardinality-independent, i.e., it applies equally to finite and infinite games, at the price of only having certain easily-identified fixed points correspond to equilibria. Equally interestingly, our differently-made choice means that our construction does not depend on games being *strategic*; this is of both technical and practical relevance, partly because the array-structure of strategic games mandates that the agents are strategic in their analysis (rather than it being optional) and partly because array-shaped games are of exponential size in the number of agents, which is typically prohibitively expensive for all but toy examples and special cases. Along with our existence guarantee in the finite case and the fact that our development closely mimics Nash’s, we consider our proposal for a more abstract and flexible notion of games, C/P games, and the alternative view on Nash equilibria they facilitate, to be a main contribution of this paper. Another main contribution is a characterisation of our style of compromises directly in the original game, in particular because the derivation of the characterising predicate from the Nash equilibrium predicate is instructive in its own right.

In [12], Nash proved that all finite strategic games have a *probabilistic* Nash equilibrium. A detailed proof using Brouwer’s Fixed Point Theorem [1] is given in [13]. More concretely, Nash observed that the set of finite strategic games can be embedded in the set of (continuous) strategic games where each agent’s set of strategies is comprised of probability distributions, in this case over the agent’s original strategies. That the latter games always have Nash equilibria follows by the existence of fixed points of a given *update* function. To illustrate what this means, let us consider a concrete strategic game. The cells of the below array are the possible outcomes of the game. The first number in a cell is the *payoff* to agent v , who chooses the row, and the second number is the payoff to agent h , who chooses the column.

$$\begin{array}{cc}
 & \begin{array}{cc} h_1 & h_2 \end{array} \\
 \begin{array}{c} v_1 \\ v_2 \end{array} & \begin{array}{|cc|} \hline 0, 1 & 1, 0 \\ \hline 1, 0 & 0, 1 \\ \hline \end{array}
 \end{array} \tag{1}$$

The Nash equilibrium predicate (and thus Nash’s update function) involves considerations of, e.g., v wanting to move from the top-left cell to the lower-left cell because the payoff is better there — we shall informally say that v is *unhappy* with the top-left cell. In turn, h is unhappy with the lower-left cell, v is unhappy lower-right, and h is unhappy top-right. More formally, this means that the game does not have a pure Nash equilibrium, i.e., a cell where all agents are happy. A probabilistic Nash equilibrium, prescribing a compromise between multiple outcomes, arises if both agents choose between their two options with equal weight, with *expected* payoffs of $1/2$ to each.

2 Nash’s Theorem

We have informally described strategic games as being arrays. Formally:

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

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

Let s range over $\mathcal{S}_{\mathcal{A}}$ and let s_a be the a -projection of s .

In terms of Nash equilibria, agents in a strategic game are free to change the entry in their dimension of the above cartesian product but must leave other (agents’) entries unchanged in their search for a better outcome.

Definition 2 Strategy s is a Nash equilibrium for strategic game G^s if¹

$$\text{Eq}_{G^s}^N(s) \quad \triangleq \quad \forall a \in \mathcal{A}, s' \in \mathcal{S}_{\mathcal{A}}. \quad (\forall a' \in \mathcal{A} . a \neq a' \Rightarrow s_{a'} = s'_{a'}) \\ \Downarrow \\ \neg(P(s, a) < P(s', a))$$

As noted, Nash equilibria do not always exist directly in a strategic game and, to have a formal notion that can be guaranteed to form compromising equilibria, Nash considered *individual probabilities* and *probability profiles*.

¹Nash’s notation for our s' is “ $s_{-a}; \sigma$ ”, i.e., s with something else (from \mathcal{S}_a) in the a -position; our contribution starts partly from the explicit quantification over s' seen here.

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

$$\begin{aligned} W^{\mathcal{S}_a} &\triangleq \{w_a : \mathcal{S}_a \rightarrow [0, 1] \mid (\sum_{\sigma \in \mathcal{S}_a} w_a(\sigma)) = 1\} \\ W^{\mathcal{S}_{\mathcal{A}}} &\triangleq \bigotimes_{a \in \mathcal{A}} W^{\mathcal{S}_a} \end{aligned}$$

The overall probability that collectively is assigned to a strategy profile by a given probability profile and its *expected-payoff* function are as follows.

Definition 4 (Expected Payoff) Given finite $\langle \mathcal{A}, \mathcal{S}_{\mathcal{A}}, P \rangle$ with associated probability profiles, $W^{\mathcal{S}_{\mathcal{A}}}$, and $w \in W^{\mathcal{S}_{\mathcal{A}}}$, $s \in \mathcal{S}_{\mathcal{A}}$, $a \in \mathcal{A}$.

$$\begin{aligned} \mu^{\mathcal{S}_{\mathcal{A}}}(w, s) &\triangleq \prod_{a \in \mathcal{A}} w_a(s_a) \\ EP_P^{\mathcal{S}_{\mathcal{A}}}(w, a) &\triangleq \sum_{s \in \mathcal{S}_{\mathcal{A}}} \mu^{\mathcal{S}_{\mathcal{A}}}(w, s) \cdot P(s, a) \end{aligned}$$

Preparing for the statement of Nash's Theorem, we have the following.

Proposition 5 (Probabilistic Strategic Games) For a finite strategic game, $G^s = \langle \mathcal{A}, \mathcal{S}_{\mathcal{A}}, P \rangle$, $W^{G^s} \triangleq \langle \mathcal{A}, W^{\mathcal{S}_{\mathcal{A}}}, EP_P^{\mathcal{S}_{\mathcal{A}}} \rangle$ is a strategic game.

We can now articulate Nash's result and sketch its proof.

Theorem 6 (Nash [12, 13]) For any finite strategic game, G^s , there exists a $w_0 \in W^{\mathcal{S}_{\mathcal{A}}}$ that is a Nash equilibrium in $W^{G^s} : \text{Eq}_{W^{G^s}}^N(w_0)$.

Proof We follow [13]. $W^{\mathcal{S}_{\mathcal{A}}}$ is the cartesian product of each agent's $W^{\mathcal{S}_a}$. Because they involve a summation to 1, each $W^{\mathcal{S}_a}$ is the standard simplex of the vector space spanned by the elements of \mathcal{S}_a taken as unit vectors. As a result, $W^{\mathcal{S}_{\mathcal{A}}}$ is a convex polytope in the vector space spanned by $\mathcal{S}_{\mathcal{A}}$, which in particular means that it is compact. More, it is possible to define a continuous function on the 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 makes a combined change. This function has a fixed point by the generalised Brouwer's Fixed Point Theorem² [1] and any such fixed point is a Nash equilibrium [13]. \square

²“A continuous function on a non-empty, compact, convex subset of a Euclidean n-space (for finite n) has a fixed point”.

If the payoffs are rational numbers, the problem of finding (actually, approximating) a probabilistic Nash equilibrium for a finite strategic game with at least two agents is in PPAD in the size of $\mathcal{S}_{\mathcal{A}}$ [15]. In fact, it is PPAD-complete [3]. PPAD is a class of problems that can be solved using so-called directed polynomial parity arguments. The main feature that makes this complexity class interesting is that it may (or may not) separate the P and NP complexity classes [7]. For our purposes, this means that it is unlikely that we will know whether there is a polynomial-time algorithm for finding (approximating) a probabilistic Nash equilibrium for some time.

Returning to Nash’s result, we can formally relate the identified equilibria back to the original game.

Definition 7 (Nash [13]) *The strategy profiles (elements in $\mathcal{S}_{\mathcal{A}}$) that are used in a probability profile, $w \in W^{\mathcal{S}_{\mathcal{A}}}$, are those given non-zero probability.*

$$\text{Use}_{W^{\mathcal{G}^{\mathcal{S}}}}(w) \triangleq \{s \in \mathcal{S}_{\mathcal{A}} \mid \mu^{\mathcal{S}_{\mathcal{A}}}(w, s) \neq 0\}$$

Given a Nash equilibrium in a probabilistic game, the associated $\text{Use}_{W^{\mathcal{G}^{\mathcal{S}}}}$ -set consists of those (pure) strategy profiles that the equilibrium prescribes a *compromise* between. A natural question is whether there is a characterising predicate for such compromises that can be expressed in terms of $\mathcal{G}^{\mathcal{S}}$. Firstly, such a predicate might conceivably have lower complexity than attempting to identify a fully probabilistic equilibrium, albeit at the price of not knowing the concrete probabilities. Secondly, it is plausible that such a predicate would provide formal insight into what principles are behind the employed compromise formation and that this, e.g., could help us target its use for modelling and analysis. (We do not know of a characterising predicate in the case of Nash’s Theorem 6 but will address the matter for our development.)

3 C/P Games

To facilitate our discrete development, we introduce an abstract game formalism called conversion/preference (C/P) games. It is based on strategic games, with explicit conversion and preference relations accounting for the views and options available to the partaking agents.

Definition 8 (C/P Games) G^{CP} are 4-tuples $\langle \mathcal{A}, \mathcal{S}, (\succ_a)_{a \in \mathcal{A}}, (\triangleleft_a)_{a \in \mathcal{A}} \rangle$:

- \mathcal{A} is a non-empty set of agents.
- \mathcal{S} is a non-empty set of synopses.³
- For $a \in \mathcal{A}$, $\succ_a \subseteq \mathcal{S} \times \mathcal{S}$ says when agent a can convert from a synopsis (the left) to another (the right).
- For $a \in \mathcal{A}$, $\triangleleft_a \subseteq \mathcal{S} \times \mathcal{S}$ says when agent a prefers a synopsis (the right) to another (the left).

Synopses are abstractions over *strategy profiles* but, as we shall see, they can also be used to denote game situations more directly.

The preference relations, \triangleleft_a , account relatively for the penalties and rewards the agents may receive in the different outcomes albeit without explicit stipulation of concrete payoffs. (In terms of strategic games, this abstraction is referred to as *strategic-form* games [14]).

In strategic(-form) games, the cartesian-product nature of the set of strategy profiles, $\mathcal{S}_{\mathcal{A}}$, determines what alternatives are available to a particular agent. In C/P games also this aspect is abstracted out in the form of the conversion relations, \succ_a , without relying on an underlying structure of \mathcal{S} . In other words, if \mathcal{S} corresponds to the game board, the \succ_a are the rules of play.

In C/P games, Nash equilibria explicitly have the reading of being synopses that no agent *can* and *want* to move away from or, said differently, if an agent can move away, that agent does not want to move away.

Definition 9 *Synopsis s is an abstract Nash equilibrium for G^{cp} , if*

$$\text{Eq}_{G^{\text{cp}}}^{\text{aN}}(s) \quad \triangleq \quad \forall a \in \mathcal{A}, s' \in \mathcal{S}. \quad s \succ_a s' \Rightarrow \neg(s \triangleleft_a s')$$

Viewing a strategic(-form) game as a C/P game sends Nash equilibria to abstract Nash equilibria and vice versa, and we shall therefore suppress the word ‘abstract’ and, in fact, we take Definition 9 as primitive.

Proposition 10 *Given a strategic game, $G^s = \langle \mathcal{A}, \mathcal{S}_{\mathcal{A}}, P \rangle$, let*

$$\begin{aligned} s \succ_a^{G^s} s' &\quad \triangleq \quad \forall a' . (a \neq a' \Rightarrow s_{a'} = s'_{a'}) \\ s \triangleleft_a^{G^s} s' &\quad \triangleq \quad P(s, a) < P(s', a) \end{aligned}$$

³The name ‘synopsis’ is inspired by the thespian meaning of ‘abstract of a play’.

$\text{cp}(G^s) \triangleq \langle \mathcal{A}, \mathcal{S}_{\mathcal{A}}, (\succ_a^{G^s})_{a \in \mathcal{A}}, (\triangleleft_a^{G^s})_{a \in \mathcal{A}} \rangle$ is a C/P game and

$$\text{Eq}_{G^s}^N(s) \Leftrightarrow \text{Eq}_{\text{cp}(G^s)}^{\text{aN}}(s)$$

We note that the result covers *all* strategic games, whether finite, countable, or continuous, meaning in particular that both the original and the derived probabilistic games considered in Theorem 6 are C/P games, and that any and all of their Nash equilibria are accounted for by Definition 9.

Having made our game formalism as abstract as we have, we can now formalise our happiness account of Nash equilibria. To that end, we first formalise unhappiness, i.e., when an agent can and wants to move.

Definition 11 *The (free) change-of-mind relation for a is $\rightarrow_a \triangleq \succ_a \cap \triangleleft_a$. Let $\rightarrow \triangleq \bigcup_{a \in \mathcal{A}} \rightarrow_a$ and let \rightarrow^* be the reflexive, transitive closure of \rightarrow :*

$$\frac{s_1 \rightarrow s_2}{s_1 \rightarrow^* s_2} \quad \frac{}{s \rightarrow^* s} \quad \frac{s_1 \rightarrow^* s_0 \quad s_0 \rightarrow^* s_2}{s_1 \rightarrow^* s_2}$$

For $C \subseteq \mathcal{S}$, let $\overset{C}{\rightarrow} \triangleq \rightarrow \cap (C \times C)$, i.e., the sub-graph of \rightarrow that spans C .

Ruling out the possibility that a synopsis is a Nash equilibrium is as simple as finding a change-of-mind step out of it. Conversely, all agents are happy if no change-of-mind steps go out of a synopsis, i.e., if the synopsis has out-degree 0 in the change-of-mind graph.

Proposition 12 $\text{Eq}_{G^{\text{cp}}}^{\text{aN}}(s) \Leftrightarrow \text{Out}_{\rightarrow}^0(s)$

Proof \rightarrow_a is the intersection of \succ_a and \triangleleft_a . □

Following this result, and for convenience, we shall often think of a C/P game simply in terms of its change-of-mind graph(s): $\langle \mathcal{S}, (\rightarrow_a)_{a \in \mathcal{A}} \rangle$. Of course, doing this is not unambiguous as multiple C/P games will have the same change-of-mind graph(s) but the insight captured in Proposition 12 is that it suffices to know \rightarrow in order to address Nash equilibria mathematically.

4 Change-of-Mind Equilibria

For conceptual and presentation reasons, we shall introduce and independently justify the compromise-characterising predicate that applies to our

discrete Nash Theorem ahead of presenting the theorem itself. Further to Proposition 12, we see that no outcome in example (1) is a Nash equilibrium because, from each of them, some agent can go somewhere else, specifically counter-clockwise. What we shall make formal in this section is that, while both agents are never happy at the same time, they are not too unhappy either, as they will keep coming back to each of the four outcomes, so to speak. The first step in the formalisation is to read Proposition 12 to say that a synopsis is a Nash equilibrium if, try as they may, no agent can change their mind about it, and vice versa for irreflexive change-of-mind relations.

Proposition 13 *Provided $\forall s. \neg(s \rightarrow s)$, it is the case that*

$$\text{Eq}_{\text{G}^{\text{cp}}}^{\text{aN}}(s) \quad \Leftrightarrow \quad (\forall s' \in \mathcal{S} . s \rightarrow^* s' \Leftrightarrow s' = s)$$

While seemingly less interesting than Proposition 12, the result suggests the following technical definition of compromises for which no agent can make *unanticipated* changes of mind.

Definition 14 *For $C \neq \emptyset$, \xrightarrow{C} is a change-of-mind equilibrium for G^{cp} if*

$$\text{Eq}_{\text{G}^{\text{cp}}}^{\text{com}}(\xrightarrow{C}) \quad \triangleq \quad \forall s \in C, s' \in \mathcal{S} . s \rightarrow^* s' \Leftrightarrow s' \in C$$


We note that change-of-mind equilibria are defined to be subgraphs, rather than just a (compromising) set of synopses in order to explicitly bring out their dynamic nature: they are equilibria in terms of “behaviours”. Informally, change-of-mind equilibria are guaranteed to be *sustainable*, i.e., situations that no agent can abandon and, seeing that different parts of an equilibrium cannot be mutually inaccessible, movement within an equilibrium cannot become “stuck” unless the equilibrium is static. We note that change-of-mind equilibria can be seen as relaxing the equality of Proposition 13 to an equivalence. Specifically, we relax it to the equivalence of what cannot be distinguished by the change-of-mind relation. (We shall make this formal in Section 6.) Reaffirming this view, and serving as a basic sanity check for the use of the change-of-mind relation in the prescribed manner, we note that Nash equilibria are exactly the static change-of-mind equilibria.

Proposition 15 $\text{Eq}_{\text{G}^{\text{cp}}}^{\text{com}}(\xrightarrow{\{s\}}) \wedge \neg(s \rightarrow s) \Leftrightarrow \text{Eq}_{\text{G}^{\text{cp}}}^{\text{aN}}(s)$

With this, we can state our main result, although we only prove it later.

Theorem 16 For any finite C/P game, G^{cp} , there exists a $C \subseteq \mathcal{S}$ such that \xrightarrow{C} is a change-of-mind equilibrium in G^{cp} : $\text{Eq}_{G^{\text{cp}}}^{\text{com}}(\xrightarrow{C})$.

To illustrate the novel dynamic equilibria captured by Definition 14, we consider the following 8 graphs, each illustrating a different view on the depicted C/P game comprising 8 synopses and 11 change-of-mind steps.



In order top-left to top-right, bottom-left to bottom-right, the graphs track stepwise changes of mind starting from the top node and with the currently considered synopses indicated by \bullet s. Nodes written \circ and edges written $--->$ are yet to be considered. In the first step, the two nodes below the top node are included for consideration via the two \rightarrow (using the left-to-right implication in Definition 14). Along with the new expectations, the second step involves the top node being abandoned, $\bullet, \dots \rightarrow$, because the other \bullet s cannot go back to it (using the right-to-left implication in Definition 14). In the fourth step, the left node gets separated out as a change-of-mind equilibrium, \otimes , onto itself, i.e., as a Nash equilibrium. In the last step, the sub-graph in the lower line is recognised as a (dynamic) change-of-mind equilibrium. The central square is not abandoned until the fifth step because of the $--->$ going into it. (One could argue that starting from the top node is somewhat arbitrary but, in fact, starting from any other node would also have ended in one or both of the above-identified equilibria. The point is that we cannot guarantee that we will always identify both equilibria. Starting in the top node above does identify both equilibria because all the other nodes

can be reached from it.) We shall return with a formal and comprehensive graph-theoretic treatment of change-of-mind equilibria in Section 6.

In terms of our earlier example, (1), we see that the whole game constitutes the only change-of-mind equilibrium. The top-left node, for example, is included because it is a better alternative for h than the top-right node.

$$\begin{array}{ccc}
 & & \begin{array}{cc} h_1 & h_2 \end{array} \\
 & & \begin{array}{|c|c|} \hline v_1 & \mathbf{0, 1} \quad \mathbf{1, 0} \\ \hline v_2 & \mathbf{1, 0} \quad \mathbf{0, 1} \\ \hline \end{array} \\
 \begin{array}{ccc} 0, 1 \leftarrow & & 1, 0 \\ \downarrow & & \uparrow \\ 1, 0 \rightarrow & & 0, 1 \end{array}
 \end{array}$$

As a compromise, this is the same as what Nash's Theorem points to (right, with non-zero probabilities indicated by boldface) and, in fact, the average and expected payoffs in the two compromises are also identical: $1/2$ to each. The compromises are differently justified, though: by change-of-mind vs payoff-driven, and in general the two need not be related, see Section 8.

5 Fixed Point Construction

In this section, we present a Tarski-style fixed-point argument for our discrete set-up. It parallels the standard Brouwer-style fixed-point argument used for probabilistic games but there are some crucial and instructive differences. In particular, the finite and infinite cases are treated uniformly. While equilibrium existence is not guaranteed from the fixed-point argument itself, the construction we employ enables a simple additional counting argument in the finite case that suffices for existence. More generally, the construction establishes a notion of approximation for change-of-mind equilibria that, e.g., opens a door to the study of (novel classes of) infinite games.

Underlying our informal notion of stepwise change-of-mind is the following discrete update function. In analogy to Nash's update function, the function takes a compromise and puts together a new compromise based on how the agents would like to improve upon the old compromise.

Definition 17 *Given G^{CP} and $C \subseteq \mathcal{S}$, let $\mathcal{U}(C) \triangleq \bigcup_{s \in C} \{s' \mid s \rightarrow^* s'\}$.*

With this, we have the following result, covering all C/P games.

Lemma 18 *Given (any) G^{cp} , \mathcal{U} has a complete lattice of fixed points.⁴*

Proof (The empty set, \emptyset , and the full set, \mathcal{S} , are trivial fixed points of \mathcal{U} .) We note that \mathcal{U} is order-preserving ($C_1 \subseteq C_2 \Rightarrow \mathcal{U}(C_1) \subseteq \mathcal{U}(C_2)$) by construction, and that it is defined on $\mathcal{P}(\mathcal{S})$, which trivially is a complete lattice when ordered by set inclusion. We are therefore done by Tarski’s Fixed Point Theorem [20]. (The fixed-point lattice is also ordered by inclusion.) \square

The complete lattice of fixed point is given as the sets of synopses to which no new synopses can be added by an agent doing a change-of-mind step, i.e., by the synopses that are on a “maximally-extended” change-of-mind path. Evidently, not all fixed points will correspond to equilibria but those that do can be nicely characterised. Further to Definition 14, they are namely the sets of synopses from which no synopsis can be removed without breaking the “maximal-extension” property. Technically speaking, the equilibria are therefore given as the least, non-empty fixed points of our update function.

Lemma 19 *Given G^{cp} , with change-of-mind relation \rightarrow .*

$$\begin{array}{c} \text{Eq}_{G^{\text{cp}}}^{\text{com}}(\overset{C}{\rightarrow}) \\ \Downarrow \\ \mathcal{U}(C) = C \wedge C \neq \emptyset \wedge (\forall C'. \emptyset \subsetneq C' \subsetneq C \Rightarrow \mathcal{U}(C') \not\subseteq C') \end{array}$$

Proof By two direct arguments. The only interesting step is from bottom to top and showing that, for any two $s_1, s_2 \in C$, we have $s_1 \rightarrow^* s_2$. We first note that \mathcal{U} is post-fixpointed: $C \subseteq \mathcal{U}(C)$, idempotent: $\mathcal{U}(\mathcal{U}(C)) = \mathcal{U}(C)$, and order-preserving: $C_1 \subseteq C_2 \Rightarrow \mathcal{U}(C_1) \subseteq \mathcal{U}(C_2)$. By order-preservation and the assumed $\mathcal{U}(C) = C$, we have $\mathcal{U}(\{s_1\}) \subseteq C$. If $\neg(s_1 \rightarrow^* s_2)$, then $s_2 \in C \setminus \mathcal{U}(\{s_1\})$, i.e., $\mathcal{U}(\{s_1\}) \subsetneq C$. By post-fixpointed-ness, $\mathcal{U}(\{s_1\})$ is non-empty and, by assumption of least-ness, we may therefore conclude $\mathcal{U}(\mathcal{U}(\{s_1\})) \subsetneq \mathcal{U}(\{s_1\})$. This contradicts idempotency, and thus $s_1 \rightarrow^* s_2$. \square

Change-of-mind equilibria are therefore *atomic*, in the sense that neither anything smaller nor anything bigger will have the same defining properties. For finite G^{cp} , a simple counting argument shows that the complete lattice of \mathcal{U} -fixed point will have least, non-empty elements, thus guaranteeing existence. (We give a different formal proof in Sections 6 and 7, establishing also the time complexity.) For the infinite case, e.g., the following unbounded change-of-mind relation will not lead to the existence of least, non-empty elements in the fixed-point lattice because all tails are fixed points.

⁴We note that complete lattices are non-empty by definition.



We leave the infinite case for future work (but see Theorem 23).

6 Graph Characterisation and Computation

We saw in Proposition 12 that Nash equilibria are exactly the nodes in the associated change-of-mind graph of a C/P game with out-degree 0. For finite games, only cycles can prevent the existence of such synopses, which in graph terminology means that we are interested in *strongly connected components*.

Definition 20 (SCC, Shrunken Graph) Consider $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$.

- The strongly connected component (SCC) of $v \in \mathcal{V}$ is

$$[v] \triangleq \{v' \mid v \rightarrow^* v' \wedge v' \rightarrow^* v\}$$

- The set of SCCs of a graph is $[\mathcal{V}] \triangleq \{[v] \mid v \in \mathcal{V}\}$.
- The shrunken graph of $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$ over $[\mathcal{V}]$ is

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

Theorem 21 ([11, 19]) The shrunken graph of a finite directed graph can be found in linear time in the sizes of \mathcal{V} , \rightarrow , i.e., with time complexity $|\mathcal{V}|^2$.

Although ambiguous, we shall write $[G^{\text{cp}}]$ for the “shrunken” version of G^{cp} , i.e., for a game whose change-of-mind relation is the shrunken graph, \curvearrowright , of G^{cp} ’s change-of-mind relation, \rightarrow . (As noted, Nash and change-of-mind equilibria depend only on the change-of-mind relation and the ambiguity therefore does not affect the status of these concepts.)

Note that $[G^{\text{cp}}]$ ’s change-of-mind relation is constructed by first intersecting conversion and preference, and then finding the shrunken graph. This means, e.g., that for the case where $s_1^a \succ_a s_1^b$ and $s_2^a \triangleleft_a s_2^b$, with $[s_1^a] = [s_2^a]$ and $[s_1^b] = [s_2^b]$, for other reasons, $[G^{\text{cp}}]$ will not necessarily have $[s_1^a] \curvearrowright [s_1^b]$.

Further to Section 4, we are now ready to prove that change-of-mind equilibria are associated with the least equivalence relation that prevents cycles: they are the SCCs with out-degree 0 in the shrunken graph, $\text{Out}_{\curvearrowright}^0(C)$.

Lemma 22 $\text{Eq}_{G^{\text{cp}}}^{\text{com}}(\xrightarrow{C}) \Leftrightarrow \text{Eq}_{[G^{\text{cp}}]}^{\text{aN}}(C)$

Proof By two direct arguments, using Proposition 12. The only interesting step is from left to right and showing that C is an $[-]$ -equivalence class. As C is non-empty, we have an $s_1 \in C$. To prove $C \subseteq [s_1]$, consider some $s_2 \in C$. By the \Leftarrow -direction of the assumed $\text{Eq}_{G^{\text{cp}}}^{\text{com}}(\xrightarrow{C})$ applied twice, we have $s_1 \rightarrow^* s_2$ and $s_2 \rightarrow^* s_1$, i.e., $s_2 \in [s_1]$, as required. Conversely, if $s_2 \in [s_1]$, we have $s_1 \rightarrow^* s_2$, and by the \Rightarrow -direction of the assumed $\text{Eq}_{G^{\text{cp}}}^{\text{com}}(\xrightarrow{C})$, $s_2 \in C$. \square

Similar to the formalism in the preceding two sections, this lemma allows us to characterise change-of-mind equilibria from a new perspective. In this case, we see that changes-of-mind will proceed in a general “forward” direction and although we may fail to make progress, we can never go back to an earlier SCC, meaning that change-of-mind equilibria have a certain *inevitable* quality to them, albeit in a weak and non-deterministic sense.

7 Technical Summary

By Lemmas 18, 19, 22, Theorem 21, we have the following master theorem.

Theorem 23 For $G^{\text{cp}} = \langle \mathcal{A}, \mathcal{S}, (\succ_a)_{a \in \mathcal{A}}, (\triangleleft_a)_{a \in \mathcal{A}} \rangle$ and non-empty $C \subseteq \mathcal{S}$, the following are equivalent.

- \xrightarrow{C} is a change-of-mind equilibrium, $\text{Eq}_{G^{\text{cp}}}^{\text{com}}(\xrightarrow{C})$.
- C is a “shrunk” Nash equilibrium, $\text{Eq}_{[G^{\text{cp}}]}^{\text{aN}}(C)$.
- C is a least non-empty fixed point of \mathcal{U} .

For finite $\langle [\mathcal{S}], \curvearrowright \rangle$ (meaning, in particular, for finite G^{cp}), such a C exists and all can be found with time complexity $|\mathcal{S}| + \sum_{a \in \mathcal{A}} |\succ_a| + \sum_{a \in \mathcal{A}} |\triangleleft_a|$, which is bounded by $|\mathcal{A}| |\mathcal{S}|^2$. If \rightarrow is given, the time complexity is $|\mathcal{S}| + |\rightarrow|$, which is bounded by $|\mathcal{S}|^2$. For infinite G^{cp} , a complete lattice of \mathcal{U} -fixed point still exists and any infinitely-descending chain of gradually smaller fixed-points amounts to gradually better approximations of a change-of-mind equilibrium.

First, we note that the theorem subsumes our main Theorem 16.

Second, we note that the theorem establishes that change-of-mind equilibria are sustainable, inevitable, and atomic in precise technical senses. As

we also saw with our example in Section 4, this means that change-of-mind equilibria enjoy a *self-stabilising* property.

Third, we stress that the latter two characterisations of C in the theorem provide a close analogue to Nash’s Theorem 6 and its proof, only these two are discrete as opposed to W^{G^s} and Nash’s update function that are continuous.

Fourth, we re-assert that it is clear additional bonuses of our development i) that we can provide also the top characterisation of the identified compromising equilibria in the theorem, ii) that it is using a predicate that is defined directly over the considered G^{cp} , iii) that the predicate formally espouses a principle of compromise formation, and iv) that the predicate closely resembles the Nash equilibrium predicate.

Fifth, the approximation property for the infinite case immediately suggests that limits, if they exist, may be interesting. Take, for example, the game where agents attempt to fold a sheet of paper to make it have the least surface area. While not an actual outcome, 0 would be a “limit” Nash equilibrium, i.e., a surface area that all agents would be happy with.

Last, we stress that the low complexity of finding change-of-mind equilibria combined with the possibly sub-exponential size of C/P games makes our formalism have practical relevance as well, as we shall discuss in Section 10.

8 Strategic Compromises

Example (1) betrays the substantial differences that may exist between the compromises prescribed by Nash’s probabilistic Theorem 6 and our discrete Theorems 16 and 23. We now show that all possible configurations can be observed when comparing the two forms of compromises (i.e., whether and how they overlap) and also that their expected/average payoffs are unrelated.

Change-of-Mind are Probabilistic Generalising example (1) 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 an equilibrium. This is illustrated below right, with six involved outcomes and average payoffs of $1/2$ to each agent.

	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

$0, 1$	←	$1, 0$
↓		↑
$1, 0$	≻	$0, 1$
	↓	↑
	$1, 0$	≻
		$0, 1$

The only probabilistic Nash equilibrium arises again if both agents choose between their options with equal probability, for expected payoffs of 1/3 to each, and involving all nine possible outcomes.

Probabilistic are Change-of-Mind A different generalisation of example (1) 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	$1, -7$

$0, 1$	←	$1, 0$	→	$0, 1$
↓		↑		↓
$1, 0$	≻	$0, 1$	←	$1, -7$
	←		←	

If v puts all weight on one row, h will want to put all weight on the columns where he gets a payoff of 1, which will make v reassign weights toward the other row. If v puts weight on both rows, h will unequivocally prefer the first column to the third, i.e., h_3 will be given probability 0. In other words, the only probabilistic Nash equilibrium involves v_1, v_2, h_1 , and h_2 with equal probabilities and expected payoffs of 1/2. The only change-of-mind equilibrium involves all six outcomes, with average payoffs of 1/2 to v and $-2/3$ to h.

Disjoint Compromises Similarly, 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

$0, 1$	←	$1, 0$
↓		↑
$1, 0$	≻	$0, 1$
	↓	↑
	$1, 0$	≻
		$0, 1$

In any probabilistic Nash equilibrium, agent v chooses strategy v_4 with full weight, and expected payoffs of 0. 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 and average payoffs of 1/2.




Figure 1: Conversions and preferences in the prisoners' dilemma — solid lines are for agent v (left payoff), dotted lines are for agent h (right payoff)

Non-Trivial Overlaps The strategic 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 and expected payoffs of $1/3$. The only change-of-mind equilibrium is as shown, with average payoffs of $1/2$.

Summary Focusing narrowly on their prescribed compromises for strategic games, we have not been able to separate change-of-mind and probabilistic Nash equilibria quantitatively, i.e., in terms of a measure. In particular, the examples above show that neither notion consistently results in smaller compromises than the other nor higher average/expected payoffs. Indeed, the two notions appear to be of independent interest and to have complementary qualitative relevance. (We shall show in Section 10 that change-of-mind equilibria appear to account nicely for steady-states in biochemical systems.)

9 Strategic vs Non-Strategic Analyses

The prisoners' dilemma is a strategic analysis of the prison time faced by two co-accused, depending on whether they remain quiet or fink on each other.



Figure 2: Agent-specific and global change-of-mind in the prisoners' dilemma

	h_{quiet}	h_{fink}
v_{quiet}	-1, -1	-3, 0
v_{fink}	0, -3	-2, -2

Figure 1 presents the information that is brought out when the prisoners' dilemma is presented in C/P form. We see that the two preference relations can be obtained from each other by flipping around the top-left/lower-right diagonal in the two right-most graphs in the figure. In particular, this means that the two agents have opposite interests, except on the diagonal, which gives rise to the asymmetry of the \prec s inside the graphs: v prefers the left outcomes to the right, while h prefers the upper-right to the lower-left. (This, in fact, accounts for the 'dilemma'.) Figure 2 presents the change-of-mind relations that result from this, with the known Nash equilibrium indicated in bold-face, in spite of both agents preferring the upper-left outcome to it.


Our main concern here is a degenerate version of the prisoners' dilemma in which we consider two agents that share two tokens that they play for: an agent with a token can take the token from the other agent 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


As before, the lower-right cell of the strategic version of blink-and-you-lose is the only Nash and change-of-mind equilibrium. By contrast, a non-strategic C/P version can provide alternative views and a more nuanced analysis, e.g., by taking as synopses the three game situations: A) agent v has both tokens; B) the agents have a token each; C) agent h has both tokens. According to the rules of the game, agent v prefers his winning situation, A, to B and C, and the neutral B to C, his losing situation. Similarly, for agent

h and C over A and B, and B over A. We now consider different types of conversion relations, corresponding to different player types, so to speak.

Foresight: An agent realises that he can win by taking the opponent’s token faster than the opponent can react, i.e., agent v can convert B to A by outpacing agent h. Agent h, in turn, can convert B to C. This version of the game has two static change-of-mind equilibria (that, thus, are also Nash equilibria): A and C.



Hindsight: An agent, say v, analyses what would happen if he does not act. In case h acts, the game would end up in C and v loses, and v therefore concludes that he could have prevented the C outcome by acting. In other words, it is within v’s power to convert C to B. Similarly for agent h from A to B. This version of the game has one static change-of-mind equilibrium (that, thus, is also a Nash equilibrium): B.



Omnisight: The agents have both hindsight and foresight, resulting in a C/P game with one change-of-mind equilibrium covering all outcomes (implying that no Nash equilibrium exists).



The C/P-game versions of blink-and-you-lose make it clear that a strategic analysis mandates subtle and potentially unintended or undesired analytic capabilities on the part of the agents. In the strategic version of blink-and-you-lose, for example, h necessarily considers the lower-right cell as an alternative to the lower-left although the latter is actually a final state of the modelled game. More to the point, actual game play can go from leave, leave to take, take but not through a winning situation the way the analysis prescribes, which could affect the appropriateness of any identified compromise/equilibrium. By contrast, the C/P-game versions can distinguish possible actions, such as in the foresight example, from what amounts to co-actions, e.g., “rectification” steps for non-actions, as in the hindsight example and presumably many more, if appropriate and as desired.

10 Systems Biology Games

Systems biology is the study of biochemical entities and how their chemical affinities (read: incentives) combine and lead to systems functionality, relative to given concentrations (read: resources). In other words, systems biology addresses the economic theory of biochemistry. (See also, e.g., [10].)

In [4], Lescanne, Vestergaard et al show that the systems-biology analyses due to Kauffman [8, 9] and Thomas [21] are, in fact, ad hoc uses of (exponential-sized array-shaped) C/P games and change-of-mind equilibria. The analyses are widely applied throughout biology, in particular in theoretical biology and for the study of model organisms and biochemical *motifs*. Originally conceived for *gene-regulation analysis*, they do a remarkably good job of identifying dynamic steady-states, e.g., among gene-expression profiles. In other words, the analyses precisely account for the dynamic notion of *homeostasis* [2] at the gene-expression level of abstraction and elsewhere. (The largest application we have heard of involves 100 variables/agents [5].)

Independently discovered ([16]), Vestergaard et al [17] proposes an alternative model construction, *Cascaded Games*, for a more general problem (namely for analysing arbitrary *influence graphs*, rather than just influence graphs that are associated with particular types of functional *updates*). Instead of constructing an exponential-sized array-shaped game containing, e.g., all possible expression profiles, Cascaded Games first identify the possible points of interaction between objects and catalysts, i.e., where *cascading* may take place. It then constructs a game around these that is typically of low-polynomial size.⁵ Little or no information seems to be lost by considering smaller games and, in some cases, the cascading approach allows us to identify multiple instances of a steady-state that can be arrived at through different *pathways* [17]. More, our agent- and incentive-centric approach means that we formally work with first-class notions of *catalysts* and, in particular, *inhibitors*, and this allows us to make substantial additional information available. The underpinning game theory is as follows.

⁵The human genome is thought to contain at least 20,000 genes, of which circa 2,000 *regulate* others. I.e., a Kauffman/Thomas model for the human genome would have $2^{20,000}+$ nodes. Using a data set for *B.Subtilis* [6] scaled non-trivially to human genome size, our Cascaded Games tool [18] constructs a game with 78,067 synopses and 266,043 change-of-mind steps, and identifies all its (pre- and sub-) equilibria in 36 hours on a laptop.

Definition 24 (Inhibitive Cooperative C/P Games [17]) iG^{CP} are 3-tuples $\langle \mathcal{A}, \mathcal{S}, \rightarrow \rangle$

- \mathcal{A} is a non-empty set of agents.
- \mathcal{S} is a non-empty set of synopses.
- $\rightarrow \subseteq \mathcal{S} \times \mathcal{S} \times \mathcal{P}(\mathcal{A}) \times \mathcal{P}(\mathcal{A})$ says when a group of agents, the third component, can (jointly) catalyse the change of one synopsis to another, while the agents from the fourth component can (jointly) inhibit this.

The implicit conversion relation in this definition will typically be the relevant stoichiometric laws of the biochemical objects being considered, while the implicit preference relations will capture chemical affinities and how these are changed by the catalysts and the inhibitors. I.e., the change-of-mind relation will amount to the chemical reactions that have non-negligible kinetics according to the model, oriented as dictated by the catalysts. As implied, our Cascaded Games tool [18] constructs inhibitive cooperative C/P games from some set of entities, \mathcal{E} , where $\mathcal{A} \subseteq \mathcal{E}$ and $\mathcal{S} \subseteq \mathcal{P}(\mathcal{E})$. When doing this, it is natural to insist that a change-of-mind step may not involve synopses that contain the inhibitors but our real interest in the inhibitors lies in Definitions 26 and 27, which capture the operational effect of inhibition.

Definition 25 ([17]) Given iG^{CP} , $+(iG^{\text{CP}})$ is the inhibition-free sub-game:


$$+(iG^{\text{CP}}) \triangleq \langle \mathcal{A}, \mathcal{S}, \{ \langle s_1, s_2, A, \emptyset \rangle \mid \langle s_1, s_2, A, \emptyset \rangle \in \rightarrow \} \rangle$$

Definition 26 (Pre- and Preventable Equilibria [17]) Given iG^{CP}

- $\text{pre-Eq}_{iG^{\text{CP}}}^{\text{com}}(\xrightarrow{C}) \triangleq \text{Out}_{+(iG^{\text{CP}})}^0(C)$
- $\text{Preventable}(\xrightarrow{C}) \triangleq \text{In}_{+(iG^{\text{CP}})}^0(C) \wedge \text{pre-Eq}_{iG^{\text{CP}}}^{\text{com}}(\xrightarrow{C})$

where In^0 stands for in-degree 0.

A change-of-mind equilibrium, Eq^{com} , is also a change-of-mind pre-equilibrium, $\text{pre-Eq}^{\text{com}}$, but not always vice versa. In particular, some pre-equilibria may have outgoing change-of-mind steps; we say they are *collapsible*. But, any outgoing steps will have a non-empty set of inhibitors, implying that collapsible pre-equilibria can be actively sustained. Preventable pre-equilibria are the dual of this. To avoid introducing biochemistry, consider the following adaptation of blink-and-you-lose with foresight (constructed using [18]), where an agent gets both tokens only if the other fails to react.



Cascaded Game : blink-and-you-lose_CG_[MinPrev]


The boxes indicate change-of-mind pre-equilibria, with solid boxes indicating change-of-mind equilibria proper. The two solid boxes are purple in colour to indicate that the equilibria are preventable. In other words, our analysis can now be extended to say that no agent can single-handedly win blink-and-you-lose, and that the neutral situation can be sustained *ad nauseam*.

Collapsible pre-equilibria require only minimal inhibition to be sustained, so to speak. To address the situation where all possible inhibitions are on, in a similar manner of speaking, we make the following definition, where we remove inhibited changes-of-mind prior to constructing the shrunken game.

Definition 27 (Sub-Equilibria [17]) *Given* iG^{cp}

- $\text{sub-Eq}_{iG^{cp}}^{\text{com}}(\xrightarrow{C}) \triangleq \text{Out}_{[+(iG^{cp})]}^0(C)$

All pre-equilibria will contain at least one sub-equilibrium. (Sub-equilibria may also occur outside a pre-equilibrium.) To illustrate, we consider the omniscient version of the above adaptation of blink-and-you-lose [18].



Cascaded Game : blink-and-you-lose_CG_[MinMaxPrev]

The (blue) dashed box indicates a change-of-mind sub-equilibrium, with the same interpretation as above, except that we identify it differently; note that no collapsible pre-equilibrium exists in this case.

As mentioned, any Eq^{com} is a $\text{pre-Eq}^{\text{com}}$ and any $\text{pre-Eq}^{\text{com}}$ contains at least one $\text{sub-Eq}^{\text{com}}$. In other words, we can generalise our main Theorem 16.

Theorem 28 ([17]) *For any finite inhibitive cooperative C/P game, iG^{cp} , there exist $C_p, C_s \subseteq \mathcal{S}$ such that $\xrightarrow{C_p}$ and $\xrightarrow{C_s}$ are change-of-mind pre- and sub-equilibria in iG^{cp} : $\text{pre-Eq}_{iG^{\text{cp}}}^{\text{com}}(\xrightarrow{C_p})$, $\text{sub-Eq}_{iG^{\text{cp}}}^{\text{com}}(\xrightarrow{C_s})$. For finite iG^{cp} , all $\text{pre-Eq}_{iG^{\text{cp}}}^{\text{com}}$ and $\text{sub-Eq}_{iG^{\text{cp}}}^{\text{com}}$ can be found with time complexity $|\mathcal{S}| + |\rightarrow|$.*

11 Conclusion

We have proved a new Nash Theorem. It’s main distinguishing features are that i) it is discrete, ii) it has low complexity for finding all equilibria, iii) it applies to strategic as well as non-strategic *simultaneous* games, i.e., to C/P games that iv) need not have either real-valued or any other explicit form of payoffs, and v) the equilibria exist in the original game via a Nash Equilibrium-like predicate that vi) makes the equilibria dynamic in nature. The simplest technical comparison between Nash’s and our development is that Nash uses Brouwer’s and we use Tarski’s Fixed Point Theorem. Additionally, we noted that our construction is cardinality independent, suggesting future work, e.g., on “limit” Nash equilibria via \mathcal{U} -approximation.

No obvious quantitative comparison favours the compromises fingered by either Nash’s or our theorem. Qualitatively, the dynamic nature of our equilibria seems to make our theorem particularly applicable to systems and life sciences. More work is needed before the view on game theory that is captured in the novel concepts of C/P games and the change-of-mind relation can be comprehensively accounted for. While Nash-style game theory typically is said to be *non-cooperative*, the work we report in Section 10 seems to suggest that our formalism facilitates both cooperative and non-cooperative analyses, as well as introducing the concept of inhibiting agents as known from biochemistry, along with the novel concepts of pre- and sub-equilibria. Our discussion in Section 9 also makes it clear that even in the absence of strategic analytic abilities on the part of the agents, not everything they may

wish to consider will directly reflect what is physically possible, i.e., change-of-mind is not just about doing or acting, as some of our terminology may imply. Indeed, and as we also saw, specific conversion relations may equip the agents with arbitrarily subtle and/or profound analytic capabilities.

Acknowledgements We thank John F. Nash, Jr. and Milan Vlach for their kind attention, and for suggestions and comments.

References

1. L. E. J. Brouwer. Über Abbildungen von Mannigfaltigkeiten. *Mathematische Annalen*, 71:97–115, 1912.
2. W. Bradford Cannon. *The Wisdom of the Body*. W. W. Norton, New York, 1932.
3. Xi Chen and Xiaotie Deng. Settling the complexity of 2-player Nash-equilibrium. *Electronic Colloquium on Computational Complexity (TR05-140)*, 2005. Appeared at the 47th Annual Symposium on Foundations of Computer Science (FOCS 2006).
4. Chafika Chettaoui, Franck Delaplace, Pierre Lescanne, Mun'delanji Vestergaard, and René Vestergaard. Rewriting game theory as a foundation for state-based models of gene regulation. In *Proceedings of the Fourth International Conference on Computational Methods in Systems Biology, Lecture Notes in Bioinformatics 4210*, 2006.
5. Reinhard Laubenbacher et al. The Discrete Visualizer of Dynamics tool. <http://dvd.vbi.vt.edu/>.
6. Kunihiko Hiraishi and Hirofumi Doi. Estimation of gene regulation using expression profiles by gene disruption and comprehensive sequence analysis on gene regulatory regions. Poster 1P-067, MBSJ Forum, 2006.
7. David S. Johnson. The NP-completeness column: Finding needles in haystacks. *ACM Transactions on Algorithms*, 3(2), 2007.
8. Stuart A. Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22:437–467, 1969.

9. Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
10. John Maynard Smith. *Evolution and the Theory of Games*. CUP, 1982.
11. Kurt Mehlhorn and Stefan Näher. *The LEDA Platform of Combinatorial and Geometric Computing*. Cambridge University Press, 1999.
12. John F. Nash, Jr. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.
13. John F. Nash, Jr. Non-cooperative games. *Annals of Mathematics*, 54, 1951.
14. Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.
15. Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994.
16. Jittisak Senachak, Mun'delanji Vestergaard, and René Vestergaard. Rewriting game theory applied to protein signalling in MAPK cascades. Research Report IS-RR-2006-007, JAIST, May 2006.
17. Jittisak Senachak, Mun'delanji Vestergaard, and René Vestergaard. Cascaded games. In *Proceedings of the Second International Conference on Algebraic Biology, Lecture Notes in Computer Science 4545*, 2007.
18. Jittisak Senachak and René Vestergaard (with Mun'delanji Vestergaard). The Cascaded Games Tool. <http://cascade.jaist.ac.jp>.
19. Robert E. Tarjan. Depth first search and linear graph algorithms. *SIAM Journal on computing*, pages 146–160, Januar 1972.
20. Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.
21. René Thomas. Boolean formalization of genetic control circuits. *Journal of Theoretical Biology*, 42(3):563–585, 1973.