





Expectation vs. Reality: Towards Verification of Psychological Games

Marta Kwiatkowska¹ , Gethin Norman^{1,2} ,
David Parker¹ , and Gabriel Santos¹ 

¹ Department of Computer Science, University of Oxford, Oxford, UK
{marta.kwiatkowska,david.parker,gabriel.santos}@cs.ox.ac.uk

² School of Computing Science, University of Glasgow, Glasgow, UK
gethin.norman@glasgow.ac.uk

Abstract. Game theory provides an effective way to model strategic interactions among rational agents. In the context of formal verification, these ideas can be used to produce guarantees on the correctness of multi-agent systems, with a diverse range of applications from computer security to autonomous driving. Psychological games (PGs) were developed as a way to model and analyse agents with belief-dependent motivations, opening up the possibility to model how human emotions can influence behaviour. In PGs, players’ utilities depend not only on what *actually* happens (which strategies players choose to adopt), but also on what the players had *expected* to happen (their belief as to the strategies that would be played). Despite receiving much attention in fields such as economics and psychology, very little consideration has been given to their applicability to problems in computer science, nor to practical algorithms and tool support. In this paper, we start to bridge that gap, proposing methods to solve PGs and implementing them within PRISM-games, a formal verification tool for stochastic games. We discuss how to model these games, highlight specific challenges for their analysis and illustrate the usefulness of our approach on several case studies, including human behaviour in traffic scenarios.

1 Introduction

Probabilistic model checking is a well established technique for formally verifying computerised systems that operate in uncertain or stochastic environments. In order to verify systems comprising multiple autonomous agents and/or those involving human interactions, various models and concepts from game theory have been adapted for probabilistic model checking. Stochastic games, in particular, have shown to be a versatile and useful formalism to model and study situations involving collaboration or competition among agents, successfully applied to, for example human-in-the-loop autonomous systems [12], robot navigation in the presence of humans [18] and attack-defence scenarios [2].

While traditional game theory is often used to model human decision making, it is unable to model situations such as emotional response and social norms,

where the utilities that players aim to maximise can depend on their *beliefs*. This inadequacy has been pointed out in [13], which proposed the seminal model of *psychological games* (PGs). In these games, a player’s utility depends not only on what *actually* happens in the game (i.e., which strategies are chosen by the players), but also on what the players had *expected* to happen (i.e., their belief as to the future behaviour of the other players).

This class of models makes it possible to consider different aspects that contribute to human decision-making, such as regret, trust, fear, reciprocity and fairness, and how these may influence player behaviour. Crucially, psychological game predictions have been reproduced in human experiments, thus supporting the notion of belief-dependent motivations [3]. This has been of particular relevance in economics, when trying to predict and understand how people behave regarding non-material payoffs. Naturally, though, as autonomous computerised systems become more commonplace, ensuring that they interact safely and efficiently with humans will also require this kind of reasoning.

In this paper, we make the first steps towards a more practical approach to modelling and analysing PGs, and in applying them to other scenarios. We begin with one-shot (normal form) games, considering the *normal form psychological games* (NFPGs) proposed in [13]. We work with the commonly employed solution concept of *Nash equilibria* (NE), which establishes rational strategies for a game to be those where no player has an incentive to unilaterally deviate from their strategy. Using the psychological extension of NE from [13], we propose an approach to finding optimal equilibria for NFPGs using support enumeration and non-linear programming, and highlight why computing equilibria for such games is more computationally challenging.

We next investigate extensive (multi-stage) games with psychological payoffs, under the assumption that beliefs are *local* and *state*-based. We do so by considering an extension of concurrent stochastic games (CSGs) whose reward functions can depend on both the actions taken and beliefs about those actions, proposing a method to find equilibria for finite-horizon cumulative rewards using backward induction. We develop prototype tool support for psychological games, building on the PRISM-games model checker for stochastic games [23]. Using this, we model and analyse a variety of psychological games, notably studying human behaviour in several different traffic scenarios, and showcase the analysis and insights made possible by our approach.

Related work. Psychological games were proposed in [13] and shown to admit standard game-theoretic techniques such as backward induction under some restrictions. However, they assume a fixed payoff structure and do not support belief inference or updating. *Dynamic psychological games* [5,4,6] address some of these limitations by allowing belief update. More specifically, they remove restrictions enforced in [13] that make beliefs endogenous to the games, and propose a *forward induction* algorithm with belief updates which allows for more sophisticated analysis. In [31], *fairness equilibria* are introduced as an extension of the framework established in [13], where the payoff of each player is defined

as a combination of a material payoff and a psychological payoff whose value depends on how fairly they think they are being treated.

From an application perspective, conventional game theory has been employed to model a range of road user behaviours [10,19,34,7,29,11], including merging into traffic and speed selection, and has been able to explain how informal norms of behaviour can develop among road users and be sustained even if these informal norms violate the formal regulations of the traffic code. In [19], the authors point out that autonomous agents should have inferable behaviour and model a pedestrian crossing interaction as a repeated bimatrix Stackelberg game in order to measure and establish a bound to inferability loss. A more complex interaction involving vehicles, pedestrians and cyclists at a crossing was investigated as a non-zero sum game in [7], which showed that real life behaviour corresponded to an equilibrium strategy that went against Norwegian traffic laws. This example served as inspiration for [29], which used Bayesian games and examined possible differences in the strategies pedestrians and cyclists would be likely to adopt when considering autonomous and human drivers. Another example is [11], which also considers a game-theoretic stochastic model to analyse interactions among pedestrians, autonomous and regular vehicles and investigates strategies for conflict resolution in uncontrolled traffic environments based on Stackelberg equilibria. However, to the best of our knowledge, psychological games have not been explored in road user scenarios.

The verification community has developed various software tools with support for Nash equilibria, such as PRALINE [8], EAGLE [35] and EVE [17], but we are aware of no tool support for psychological equilibria computation, in either normal or extensive forms. For probabilistic systems, model checkers such as PRISM [22] and Storm [16] support a wide range of probabilistic models, with partially observable Markov decision processes providing an alternative way to reason about belief, over (unobservable) states rather than strategies. PRISM-games [23] provides verification and equilibria synthesis for various types of stochastic games including CSGs, but until now not for psychological variants.

2 Preliminaries

We first recall normal form games (NFGs), over which we define Nash equilibria (NE), and then proceed by defining the psychological equivalents: normal form psychological games (NFPGs) and psychological Nash equilibria (PE).

Classical games. We will write $Dist(X)$ for the set of probability distributions over a finite set X .

Definition 1 (Normal form game). A (finite, n -person) normal form game (NFG) is a tuple $N = (N, A, u)$ where:

- $N = \{1, \dots, n\}$ is a finite set of players;
- $A = A_1 \times \dots \times A_n$, A_i is a finite set of actions available to player $i \in N$ and $A_i \cap A_j = \emptyset$ for $i \neq j \in N$;

– $u = (u_1, \dots, u_n)$ and $u_i: A \rightarrow \mathbb{Q}$ is a utility function for player $i \in N$.

In an NFG, each player $i \in N$ simultaneously chooses an action $a_i \in A_i$ from their action set and each player j then receives utility $u_j(a_1, \dots, a_n)$. Two-player NFGs are often called *bimatrix games* since they can be represented by two matrices $Z_1, Z_2 \in \mathbb{Q}^{l \times m}$ where $A_1 = \{a_1, \dots, a_l\}$, $A_2 = \{b_1, \dots, b_m\}$, $Z_1(i, j) = u_1(a_i, b_j)$ and $Z_2(i, j) = u_2(a_i, b_j)$. Below, we assume a fixed NFG $\mathbf{N} = (N, A, u)$.

Definition 2 (Strategy and strategy profile). A (mixed) strategy for player i of NFG \mathbf{N} is a distribution $\sigma_i \in \text{Dist}(A_i)$, specifying the probability of choosing each action in its action set. A strategy profile of \mathbf{N} (or just profile) is a tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of strategies for all players.

The *expected utility* of player i under strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is:

$$u_i(\sigma) \stackrel{\text{def}}{=} \sum_{(a_1, \dots, a_n) \in A} u_i(a_1, \dots, a_n) \cdot \left(\prod_{j=1}^n \sigma_j(a_j) \right).$$

We let $\Sigma_{\mathbf{N}}^i = \text{Dist}(A_i)$ denote the set of all player i strategies, $\Sigma_{\mathbf{N}} = \prod_{i \in N} \Sigma_{\mathbf{N}}^i$ the set of all strategy profiles and $\Sigma_{\mathbf{N}}^{-i} = \prod_{j \neq i} \Sigma_{\mathbf{N}}^j$ the set of strategy tuples for all players except i . For strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and player i strategy σ'_i , we define the strategy tuple $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ and strategy profile $\sigma_{-i}[\sigma'_i] = (\sigma_1, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_n)$.

Definition 3 (Support). The support $Q_i \subseteq A_i$ of a strategy σ_i for player i is the set of actions it chooses with positive probability, i.e., $Q_i = \{a_i \in A_i \mid \sigma_i(a_i) > 0\}$. The support of a profile σ is the product of the supports of its individual strategies σ_i .

We now define the notion of *Nash equilibria* (NE), which are strategy profiles for which there is no incentive for any player to unilaterally change their strategy.

Definition 4 (Best response). For strategy tuple $\sigma_{-i} \in \Sigma_{\mathbf{N}}^{-i}$, a best response to σ_{-i} for player i is a strategy $\sigma_i^* \in \Sigma_{\mathbf{N}}^i$ such that $u_i(\sigma_{-i}[\sigma_i^*]) \geq u_i(\sigma_{-i}[\sigma_i])$ for all $\sigma_i \in \Sigma_{\mathbf{N}}^i$.

Definition 5 (Nash equilibrium). A strategy profile σ^* is a Nash equilibrium (NE) and $\langle u_i(\sigma^*) \rangle_{i \in N}$ NE values if σ_i^* is a best response to σ_{-i}^* for all $i \in N$.

Since multiple NE can exist for an NFG, we are also interested in finding the *optimal* equilibrium for a given criterion. In this paper, we focus on *social welfare* NE, which are those that maximise the sum of the players' utilities.

Definition 6 (Social welfare NE). An NE σ^* is a social welfare optimal NE (SWNE) and $\langle u_i(\sigma^*) \rangle_{i \in N}$ corresponding SWNE values if $u_1(\sigma^*) + \dots + u_n(\sigma^*) \geq u_1(\sigma) + \dots + u_n(\sigma)$ for all NE σ .

We are now ready to discuss *normal form psychological games* (NFPGs) [13], a generalisation of NFGs in which a player's utility can depend not only on the game's outcome (actions taken), but also on the player's *belief* as to the outcome.

The notions of actions, strategies and strategy profiles remain the same as for NFGs and, for NFPG \mathbf{N}_P , we use the same notation $\Sigma_{\mathbf{N}_P}^i$, $\Sigma_{\mathbf{N}_P}$ and $\Sigma_{\mathbf{N}_P}^{-i}$.

Beliefs and coherence. The *first-order beliefs* for player i represent their belief as to the (mixed) strategies that will be taken by the other players. So, the set of all first-order beliefs for player i is defined as $B_i^1 = \text{Dist}(\Sigma_{\mathbf{N}_P}^{-i})$. *Higher-order beliefs* are beliefs about the beliefs of other players. We will denote the k th order beliefs for player i by B_i^k and write $B_{-i}^k = \prod_{j \neq i} B_j^k$ for the set of k th order beliefs for all players other than i . Higher-order beliefs are then defined inductively by:

$$B_i^{k+1} \stackrel{\text{def}}{=} \text{Dist}(\Sigma_{\mathbf{N}_P}^{-i} \times B_{-i}^1 \times \cdots \times B_{-i}^k).$$

Notice that information about beliefs appears multiple times. This allows for correlation between different orders of belief, e.g., second-order beliefs B_i^2 assign probabilities to combinations of the strategies $\Sigma_{\mathbf{N}_P}^{-i}$ and first-order beliefs B_{-i}^1 of the other players. As in [13], we will assume that beliefs are *coherent*, meaning that this information is consistent. For example, the marginal of player i 's second-order beliefs with respect to $\Sigma_{\mathbf{N}_P}^{-i}$ should coincide with i 's first-order beliefs. The same condition is applied inductively to higher-order belief sets.

Furthermore, since players are rational and know that other players are also rational, coherency is assumed to be common knowledge and we will require beliefs to be *collectively coherent*. In other words, each player i only ever believes that another player j 's beliefs are coherent, that player j believes other players to be coherent, and so on. We will write B_i for the set for all collectively coherent higher-order beliefs for player i , and define $B = \prod_{i \in N} B_i$ to be the set of collectively coherent *belief profiles*, i.e., the set of beliefs for all players.

Psychological games. We can now formally define the psychological variant of normal form games, where the key difference is that the utility for player i now also depends on their (collectively coherent) belief $b_i \in B_i$ about the other players, as well as the actions they actually take.

Definition 7 (Normal form psychological game). A (finite, n -person) normal form psychological game (NFPG) is a tuple $\mathbf{N}_P = (N, A, u)$ where:

- $N = \{1, \dots, n\}$ is a finite set of players;
- A_i is a finite set of actions available to player $i \in N$ and $A_i \cap A_j = \emptyset$ for $i \neq j \in N$;
- $u = (u_1, \dots, u_n)$ and $u_i: (B_i \times A) \rightarrow \mathbb{Q}$ is a utility function for player $i \in N$.

As for NFGs, we can define the *expected* utility of player i for a given strategy profile σ . However, here we must now also include beliefs. More precisely, for belief $b_i \in B_i$ for player i and strategy profile σ , we write $u_i(b_i, \sigma)$ for the expected utility of player i under b_i and σ .

We can now define the notion of *psychological Nash equilibrium* (PE). While an NE for an NFG \mathbf{N} is a strategy profile $\sigma \in \Sigma_{\mathbf{N}}$, a PE for an NFPG \mathbf{N}_P is a pair $(b, \sigma) \in B \times \Sigma_{\mathbf{N}_P}$ comprising a belief profile b and a strategy profile σ . Crucially, as explained in [13], it is assumed that, in equilibrium, the beliefs b of

the players match a commonly held view of reality. In other words, each player i believes, with probability 1, that each other player j follows strategy σ_j , that player j 's beliefs match σ_{-j} , and so on. For strategy profile σ , this matching belief profile for all n players is denoted $\beta(\sigma)$.

Definition 8 (Psychological Nash equilibrium). A pair (b^*, σ^*) of belief profile $b^* = (b_1^*, \dots, b_n^*) \in B$ and strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma_{\text{NPG}}$ for NFG NPG is a psychological Nash equilibrium (PE) if:

$$b^* = \beta(\sigma^*) \tag{1}$$

$$u_i(b_i^*, \sigma^*) \geq u_i(b_i^*, \sigma_{-i}[\sigma_i]) \text{ for all } \sigma_i \in \Sigma_{\text{NPG}}^i \text{ and } i \in N. \tag{2}$$

The first condition (1) implies that, as discussed above, the players' beliefs match a commonly held view of reality. The second condition (2) matches the corresponding requirements for NEs of NFGs (see Definitions 4 and 5). As for NEs, we will generally aim to find a PE that is *social welfare optimal*, where the sum of the players' utilities is maximised.

Defining utility functions. Since we focus on psychological Nash equilibria, condition (1) above, combined with the assumption of collective coherence for higher-order beliefs, allows us to adopt a simpler formulation of an NFG's utility functions in practice. Although player i 's utility function u_i depends on its (collectively coherent, higher-order) beliefs about the other players, since we know that there is a common belief in equilibrium we can simply define u_i in terms of the strategies alone, that is, as a function of the probabilities that each player j takes each of its actions. Additionally, for simplicity, we allow each player's utility to be defined in terms of *all* the player's strategies, including the players' *own* choices of actions.

Example 1. To illustrate NFGs, let us consider the *confidence* game from [13], which comprises three players. Player 1 submits a proposal, which is randomly assigned with equal probability to player 2 or 3. They can then chose to either *accept* or *reject* this proposal. We abbreviate these actions to a_i and r_i , respectively, for player $i = 2, 3$. Player 1 has no actions to take.

The game has belief-dependent utilities, involving both first-order and second-order beliefs. We will write p_a for the *probability* that a player i chooses action a in their (mixed) strategy, $\overline{p_a}$ for the expectation of another player $j \neq i$ as to the probability p_a , i.e., the first-order belief for player j , and $\overline{\overline{p_a}}$ for the expected value of $\overline{p_a}$ from the perspective of another player, i.e., the second-order belief.

Player 1 wants the proposal to be accepted, but their satisfaction about acceptance is influenced by their belief about how likely this is to happen: being more *optimistic* means they are happier about an acceptance, but also much unhappier in the case of a rejection. Player 2 is influenced by how *confident* player 1 is about acceptance, and is more likely to accept the proposal if they believe player 1 is more confident. Player 3 always prefers rejection.

These notions are encoded in the players' utility functions as follows, where we follow [13] but show extra details of the derivation. Let us denote the probability that player 1's proposal is accepted as p_{acc} . We have $p_{acc} = (1/2) \cdot (p_{a_2} + p_{a_3})$

since players 2 and 3 are assigned the proposal with equal probability. Similarly, $\overline{p_{acc}} = (1/2) \cdot (\overline{p_{a_2}} + \overline{p_{a_3}})$ is player 1's belief as to the likelihood of acceptance, and $\overline{\overline{p_{acc}}} = (1/2) \cdot (\overline{\overline{p_{a_2}}} + \overline{\overline{p_{a_3}}})$ equals player 2's belief about $\overline{p_{acc}}$.

Player 1 has a utility of 1 in case of acceptance, plus a further utility based on their degree of optimism, i.e., in terms of belief of acceptance, of $2 \cdot \overline{p_{acc}}$. Conversely, if rejected, their utility is $-8 \cdot \overline{p_{acc}}$. Player 2 prefers to accept when $\overline{p_{acc}} > \frac{1}{6}$ so, when assigned the proposal, receives a utility of $6 \cdot \overline{\overline{p_{acc}}}$ for accepting and 1 for rejecting. Player 3 has a utility of 1 if they are assigned the proposal and reject it, otherwise 0.

Recall from above that, in equilibrium, players' beliefs must match a shared view of reality, so $p_{a_i} = \overline{p_{a_i}} = \overline{\overline{p_{a_i}}}$ for $i = 2, 3$. We therefore express the utility functions for players as expressions in terms of just p_{a_i} . Since player 1 does not choose an action and players 2 and 3 each can choose between two actions, we can write player i 's utility function as a 2×2 matrix Z_i . For each pair of actions of players 2 and 3, the value combines the utility arising when each of player 2 and 3 are assigned the proposal, weighted by probability 1/2:

$$Z_1 = \begin{matrix} & a_3 & r_3 \\ a_2 & \begin{pmatrix} 1 + p_{a_2} + p_{a_3} & 1/2 - (3/2) \cdot (p_{a_2} + p_{a_3}) \\ 1/2 - (3/2) \cdot (p_{a_2} + p_{a_3}) & -4 \cdot (p_{a_2} + p_{a_3}) \end{pmatrix} \\ r_2 & \end{matrix}$$

$$Z_2 = \begin{matrix} & a_3 & r_3 \\ a_2 & \begin{pmatrix} (3/2) \cdot (p_{a_2} + p_{a_3}) & (3/2) \cdot (p_{a_2} + p_{a_3}) \\ 1/2 & 1/2 \end{pmatrix} \\ r_2 & \end{matrix} \quad Z_3 = \begin{matrix} & a_3 & r_3 \\ a_2 & \begin{pmatrix} 0 & 1/2 \\ 0 & 1/2 \end{pmatrix} \\ r_2 & \end{matrix}$$

When determining PE for the confidence game, we can ignore player 1 since it has no actions. In an equilibrium, suboptimal actions cannot be played with positive probability. We can therefore compute a solution by encoding the problem with the following set of constraints:

$$((3/2) \cdot ((p_{a_2} + p_{a_3}) \cdot p_{a_3} + (p_{a_2} + p_{a_3}) \cdot p_{r_3}) \geq (1/2) \cdot (p_{a_3} + p_{r_3})) \vee (p_{a_2} = 0) \quad (3)$$

$$((1/2) \cdot (p_{a_3} + p_{r_3}) \geq (3/2) \cdot ((p_{a_2} + p_{a_3}) \cdot p_{a_3} + (p_{a_2} + p_{a_3}) \cdot p_{r_3})) \vee (p_{r_2} = 0) \quad (4)$$

$$(0 \geq (1/2) \cdot (p_{a_2} + p_{r_2})) \vee (p_{a_3} = 0) \quad (5)$$

$$((1/2) \cdot (p_{a_2} + p_{r_2}) \geq 0) \vee (p_{r_3} = 0) \quad (6)$$

For example, (3) must hold because either the action a_2 is optimal for player 2, i.e., the utility obtained by player 2 when action a_2 is chosen is greater than or equal to that when action r_2 is chosen under the optimal strategy of player 3, or the action a_2 is chosen with probability 0. Any assignment that satisfies all four constraints is an equilibrium. Given that $p_{a_2} + p_{r_2} = 1$, the first clause of (5) cannot be satisfied and thus we must have $p_{a_3} = 0$. This is consistent with the fact that a_3 is dominated for player 3, i.e., action r_3 always yields higher utility than a_3 . If $p_{a_3} = 0$, we have $p_{r_3} = 1$, which means that the second clause of (6) has to be false. The first clause of (6) is trivially satisfied. Constraints (3) and (4) can then be reduced to:

$$(3 \cdot p_{a_2} \geq 1) \vee (p_{a_2} = 0) \quad (7)$$

$$(1 \geq 3 \cdot p_{a_2}) \vee (p_{r_2} = 0) \quad (8)$$

We then obtain satisfying assignments by setting $p_{a_2} = 1/3$ and $p_{r_2} = 2/3$, $p_{a_2} = 1$ and $p_{r_2} = 0$ or $p_{a_2} = 0$ and $p_{r_2} = 1$ with utility vectors $u = (u_1, u_2, u_3)$ equal to $(-8/9, 1/2, 1/2)$, $(-1, 3/2, 1/2)$ and $(0, 1/2, 1/2)$, respectively. The proposal has the highest chance of being accepted in the second equilibrium, when player 1 is most confident. However, as player 3 is certain to reject, that is also the worst equilibrium for player 1, who is bound to be disappointed. The last two equilibria are social welfare optimal with a combined utility of 1. ■

3 Equilibria Computation for Psychological Games

We now propose methods for analysing NFPGs in order to determine their PEs and corresponding values. The approach builds upon techniques for the non-psychological setting, i.e., finding NEs for NFGs. For the case of two-player NFGs (bimatrix games), we can use well known approaches such as the Lemke-Howson [28] algorithm or mixed-integer programming based on regret minimisation [32]. For NFGs with more than two players, algorithms include the Govindan-Wilson [14] or Simplicial Subdivision [27], as well as search methods based on *support enumeration* [30].

We take the support enumeration approach for NFPGs, by adapting the method of [24], which has been used to find social welfare optimal NEs for NFGs in a similar fashion. This approach exhaustively inspects sub-regions of the strategy profile space, based on the idea that searching for NEs within a specific *support* (see Definition 3) of a strategy profile is computationally easier. It relies on encoding the computation of a (social welfare optimal) NE as a *non-linear programming* (NLP) problem.

The NLP encoding leverages conditions for a strategy profile to characterise an NE, presented as a lemma in [24], and based on the notion of *feasibility program* introduced in [9,30]. The lemma states that a strategy profile of an NFG is an NE if and only if any player switching to a single action in the support of the profile yields the same utility for the player, and switching to an action outside the support can only decrease its utility. We adapt that lemma here to characterise a PE of an NFPG. This result follows directly from Definition 8 and, in particular, the fact that in equilibrium the belief profile needs to correspond to the strategies being played.

Lemma 1. *A pair (b, σ) comprising a belief profile b and a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ of $\mathbf{N}_P = (N, A, u)$ is a PE if and only if (1) and the following conditions are satisfied:*

$$\forall i \in N. \forall a_i \in A_i. \sigma_i(a_i) > 0 \rightarrow u_i(b, \sigma_{-i}[\eta_{a_i}]) = u_i(b, \sigma) \quad (9)$$

$$\forall i \in N. \forall a_i \in A_i. \sigma_i(a_i) = 0 \rightarrow u_i(b, \sigma_{-i}[\eta_{a_i}]) \leq u_i(b, \sigma) \quad (10)$$

where η_{a_i} is the pure strategy that picks action a_i with probability 1.

We now extend the NLP encoding of the computation of an NE presented in [24] to the case of an NFPG $\mathbf{N}_P = (N, A, u)$. Since this encoding uses a support enumeration approach, we need to determine the social welfare optimal PE amongst strategy profiles from a fixed support $Q = Q_1 \times \dots \times Q_n \subseteq A$, i.e., for a given set of actions of each player. We first choose a *pivot* action $q_i^p \in Q_i$, which can be any action in Q_i , for each player i . The problem is then to maximise:

$$\sum_{i \in N} \left(\sum_{q \in Q} u_i(b^*, q) \cdot \left(\prod_{j \in N} p_{q_j} \right) \right) \quad (11)$$

subject to:

$$\sum_{c \in Q_{-i}(q_i^p)} u_i(b^*, c) \cdot \left(\prod_{j \in N_{-i}} p_{c_j} \right) - \sum_{c \in Q_{-i}(a_i)} u_i(b^*, c) \cdot \left(\prod_{j \in N_{-i}} p_{c_j} \right) = 0 \quad (12)$$

$$\sum_{c \in Q_{-i}(q_i^p)} u_i(b^*, c) \cdot \left(\prod_{j \in N_{-i}} p_{c_j} \right) - \sum_{c \in Q_{-i}(a_i)} u_i(b^*, c) \cdot \left(\prod_{j \in N_{-i}} p_{c_j} \right) \geq 0 \quad (13)$$

$$\sum_{q_i \in Q_i} p_{q_i} = 1 \quad \text{and} \quad p_{q_i} > 0 \quad (14)$$

for all $i \in N$, $q_i \in Q_i \setminus \{q_i^p\}$ and $a_i \in A_i \setminus Q_i$ where $Q_{-i}(c_i) = Q_1 \times \dots \times Q_{i-1} \times \{c_i\} \times Q_{i+1} \times \dots \times Q_n$, $N_{-i} = N \setminus \{i\}$ and $b^* \in B$.

The variables p_{q_i} represent the probabilities of players choosing different actions, i.e. the probability player i selects action $q_i \in Q_i$. If a satisfying assignment is found, we have a social welfare optimal PE given by the belief and strategy profiles pair (b^*, σ^*) , where, for $a_i \in A_i$, $\sigma_i^*(a_i) = p_{q_i}$ if $a_i = q_i$ and $q_i \in Q_i$, and 0 otherwise. Following condition (1) of Definition 8, we have $b^* = \beta(\sigma^*)$.

Constraints (12) and (13) enforce that the solution corresponds to a PE, encoding constraints (9) and (10), respectively, of Lemma 1 when restricting to pivot actions. This restriction is sufficient as (9) requires all actions in the support to yield the same utility. The objective function in (11) corresponds to the sum of the individual utilities of the players when they play according to the profile corresponding to the solution. By maximising it, we require the solution to be social welfare optimal. As it is possible to have more than one equilibrium for which the sum of utilities is optimal, we specify additional *lower priority* objectives to maximise individual payoffs following an increasing sequence of indices i , and thus output a payoff vector with a consistent ordering.

Example 2. Consider the two-player NFPG whose utility functions are given by the matrices in Figure 1 (left). Player 1 has no choice (we write $A_1 = \{\perp\}$) and player 2 chooses an action from $A_2 = \{a_2, b_2\}$. Player 2 is indifferent (their utility is always 0), whereas player 1's utility depends on (their expectation about) the probability p_{a_2} of a_2 being played. Therefore, player 1's expected utility (which is also the total expected utility) is a function depending only on p_{a_2} :

$$u_1(p_{a_2}) = -\frac{400}{81} \cdot p_{a_2}^2 + \frac{40}{9} \cdot p_{a_2}$$

Since the only player with a choice is indifferent between their own actions, the constraints in (12) and (13) are trivially satisfied for all supports as long as (14) is

$$Z_1 = \perp \begin{pmatrix} a_2 & b_2 \\ -\frac{400}{81} \cdot p_{a_2} + \frac{40}{9} & 0 \end{pmatrix}$$

$$Z_2 = \perp \begin{pmatrix} a_2 & b_2 \\ 0 & 0 \end{pmatrix}$$




Fig. 1: Player utilities (left) and total expected utility (right) for Example 2.

also satisfied, and thus any strategy profile (with an accompanying belief profile) is an equilibrium. Figure 1 (right) plots the total expected utility. This shows that, in order to achieve the maximum value of 1, player 2 has to randomise, picking actions a_2 and b_2 with probabilities 0.45 and 0.55, respectively. ■

The above example illustrates a contrast with (non-psychological) NFGs, and gives an indication of why computing optimal equilibria for NFPGs is more computationally challenging. For NFGs, it suffices to consider pure strategies only when finding optimal equilibria with a single active player. This can be exploited [25], avoiding the need to solve an optimisation problem for a given support. For NFPGs, the non-linearity in the function for expected utility means that pure strategies no longer suffice for an indifferent player.

4 Psychological Concurrent Stochastic Games

We next consider concurrent stochastic games (CSGs) [33], which are multi-stage games played over graphs where, at each state, players make simultaneous choices that cause the game’s state to be probabilistically updated. We present a *psychological* variant of CSGs, in which, similarly to NFPGs, a player’s utility (reward accumulated over a finite horizon) can depend on its *belief* as to the strategies to be played as well as the actions that players select. We outline a procedure to compute equilibria for a class of such games, which restricts the nature of the players’ beliefs. As in previous work for CSGs [23], we consider *subgame perfect* equilibria, which are equilibria at every state of a CSG.

Definition 9 (Concurrent stochastic game). A concurrent stochastic multi-player game (CSG) is a tuple $C = (N, S, \bar{s}, A, \Delta, \delta)$ where:

- $N = \{1, \dots, n\}$ is a finite set of players;
- S is a finite set of states and $\bar{s} \in S$ is the initial state;
- $A = (A_1 \cup \{\perp\}) \times \dots \times (A_n \cup \{\perp\})$ where A_i is a finite set of actions available to player $i \in N$, $A_i \cap A_j = \emptyset$ for $i \neq j \in N$ and \perp is an idle action disjoint from the set $\cup_{i=1}^n A_i$;
- $\Delta: S \rightarrow 2^{\cup_{i=1}^n A_i}$ is an action assignment function;
- $\delta: (S \times A) \rightarrow \text{Dist}(S)$ is a probabilistic transition function.

We assume without loss of generality that, for any $i, j \in N$, $i \neq j$, $A_i \cap A_j = \emptyset$. A given CSG C starts in the initial state \bar{s} and, when in state s , each player $i \in N$ selects an action from its available actions $A_i(s) \stackrel{\text{def}}{=} \Delta(s) \cap A_i$ if this set is non-empty, and from $\{\perp\}$ otherwise. Assuming each player i selects action a_i , the next state of the game is determined according to the distribution $\delta(s, (a_1, \dots, a_n))$. A *path* of C is a sequence $\pi = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$ where $s_k \in S$, $\alpha_k = (a_1^k, \dots, a_n^k) \in A$, $a_i^k \in A_i(s_k)$ for $i \in N$ and $\delta(s_k, \alpha_k)(s_{k+1}) > 0$ for all $k \geq 0$. We denote by $\pi(i)$ the $(i+1)$ th state of π , $\pi[i]$ the action associated with the $(i+1)$ th transition and, if π is finite, $\text{last}(\pi)$ the final state. Let $FPaths$ and $IPaths$ denote the sets of finite and infinite paths that start in the initial state, respectively.

A *strategy* for a player in C resolves the player's choices in each state. These choices can depend on the execution history and can be randomised, i.e., are of the form $\sigma_i : FPaths \rightarrow \text{Dist}(A_i)$ such that if $\sigma_i(\pi)(a_i) > 0$, then $a_i \in A_i(\text{last}(\pi))$. As for NFGs, a *strategy profile* for C is a tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of strategies for all players. For a given strategy profile σ , a probability measure Prob_C^σ over the infinite paths of C can then be defined in the standard way [20].

Psychological CSGs. In order to introduce a psychological variant of CSGs, we incorporate a notion of beliefs, and then use them to define rewards. Let B_i^A denote, as defined in Section 2, the set of (collectively coherent, higher-order) beliefs for player i , where first-order beliefs are over the set of actions A . A *belief* b_i for player i is of the form $b_i : S \rightarrow B_i^A$. It is *state-based*, in that it provides a separate belief for each state s of the CSG, and *local*, in that these beliefs give the player's expectations regarding the actions to be played in s , not about a more global notion of the player's strategy. A *belief profile* is a tuple $b = (b_1, \dots, b_n)$.

A *reward structure* for player i takes the form $r_i = (r_i^A, r_i^S)$, where $r_i^A : (S \times B_i^A \times A) \rightarrow \mathbb{Q}$ is an action reward function (which maps a state, belief and action tuple to a rational value that is accumulated when the action tuple is selected in the state, assuming a given local belief for player i in that state) and $r_i^S : S \rightarrow \mathbb{Q}$ is a state reward function (which maps each state to a rational value that is accumulated when the state is passed through).

The *utility* (or *objective*) for a player i in CSG C can be defined by a random variable $X_i : IPaths \rightarrow \mathbb{R}$ mapping infinite paths to reals. We denote by $\mathbb{E}_C^\sigma(X_i)$ the expected value of player i 's utility under σ , with respect to the probability measure Prob_C^σ . Given utilities X_1, \dots, X_n for all the players of C , we can then define (social welfare) psychological Nash equilibria, as for NFGs. We will restrict our attention to utilities that correspond to *finite-horizon* objectives, which may be used to investigate, for instance, the expected reward accumulated over k steps. Such utilities can be expressed by a finite bound $k \in \mathbb{N}$ and reward structure $r_i = (r_i^A, r_i^S)$, with corresponding random variable:

$$X_i(\pi) = \sum_{j=0}^{k-1} (r_i^A(\pi(j), b_i(\pi(j)), \pi[j]) + r_i^S(\pi(j))).$$

Psychological equilibria computation. In both [13] and [5], the authors point out the limitations of applying backward induction to computing equilibria for extensive (multi-stage) psychological games and show why that approach

cannot be applied to the general case. While a full discussion is outside the scope of this paper, it suffices to imagine a game in which a player’s utility in a given state depends on the beliefs they have about actions performed in the *preceding* state. A backward induction algorithm *can*, however, be applied under the assumption that psychological utilities at any given state have to be over *local* strategies, that is, concerning the actions taken at that same state and coherent with an equilibrium solution for the NFPG in that state. We note that a more general approach has been proposed in [5], which we leave as future work.

Using the above restriction, we devise a backward induction algorithm that builds, at each iteration, an NFPG for every state s in C according to the reward structure in s and the values computed for its successors in the previous iteration (initially 0 for all states). We then compute equilibria values and strategies by solving the corresponding NLPs following the definitions in Section 3. We compute equilibria which are *locally* social welfare optimal. Other criteria, e.g., *social cost* [25] or *social fair* [26] equilibria, which minimise the overall sum or the difference between the highest and lowest utilities, respectively, could also be applied. In the latter case, however, additional constraints would have to be added to the NLP in Section 3, which would significantly increase the complexity of the problem due to the need to numerically encode logical implications. Finding all equilibria for a CSG is generally intractable, as the number may be exponential even with respect to the size of the normal form game at each state.

5 Case Studies and Experimental Results

We have built a prototype implementation to model and solve psychological games, and used it to investigate the applicability and performance of our approach on a selection of normal form and multi-stage psychological games. We first consider two-player instances of the *ultimatum* and *reciprocity* games of [6], which exemplify how psychological games can also be used in the computation of *fairness equilibria*, as well as how psychological utilities can influence the strategies of the players. We then present two- and multi-player normal form games modelling *traffic* interactions between pedestrians, cyclists and vehicles, one of which is then extended to a psychological CSG, used to investigate how information on past decisions can influence players’ strategies.

Implementation. We build on top of the PRISM-games model checker [23], extending its existing modelling language for CSGs (in which normal form games can also be encoded as simple instances). The key difference is that in PG models the specification of reward structures needs to incorporate a player’s beliefs about the other players’ strategies. Since we currently only allow for beliefs in CSGs over local strategies (see Section 4), rewards for a state can only make reference to (the probability of) actions played in that state. Figure 2 shows a reward structure definition in our extension of the PRISM-game modelling language for the ultimatum game example (see Section 5.1, below). For simplicity, in this syntax, we just use the name of the action to denote the probability of choosing it, e.g., `reject` denotes what we refer to elsewhere as p_{reject} .

```

1 // Constants
2 const double theta1; // Player 1's reciprocity sensitivity
3 const double theta2; // Player 2's reciprocity sensitivity
4 // Rewards for player 1
5 rewards "r1"
6     [fair,reject] true : 5+theta1*(-4.5)*(2+0.5*reject);
7     [fair,accept] true : 5+theta1*(4.5)*(2+0.5*reject);
8     [greedy,reject] true : 0+theta1*(-4.5)*(-2-0.5*reject);
9     [greedy,accept] true : 9+theta1*(4.5)*(-2-0.5*reject);
10 endrewards
11 // Rewards for player 2
12 rewards "r2"
13     [fair,reject] true : 5+theta2*(-4.5)*(2+0.5*reject);
14     [fair,accept] true : 5+theta2*(4.5)*(2+0.5*reject);
15     [greedy,reject] true : 0+theta2*(-4.5)*(-2-0.5*reject);
16     [greedy,accept] true : 1+theta2*(4.5)*(-2-0.5*reject);
17 endrewards

```

Fig. 2: Reward structures for the ultimatum game, modelled in PRISM-games.

As for regular CSGs, our extension of PRISM-games constructs and stores PG models using the tool’s Java-based ‘explicit’ engine. In contrast to CSGs, reward structures for PG models are represented by symbolic expressions over variables representing action choice probabilities and cannot be evaluated prior to model checking. We use Gurobi [15] to solve the NLPs described in Section 3 for finding equilibria values of NFPGs at each state.

Efficiency and Scalability. Computing equilibria values and strategies can be a complex task, even for the simpler case of finding an arbitrary (non-optimal) equilibrium of a two-player normal form game. Finding optimal equilibria of multi-player games is considerably harder, given the increased number of supports and the non-linearity of the constraints. The addition of psychological utilities complicates the computation further, as there are no natural restrictions on how these utilities may vary given the players’ strategies and beliefs. At the state level, when looking for an optimal equilibrium, we are required to solve an NLP for each support. Given the total number of supports is exponential in the number of actions, i.e., equals $\prod_{i=1}^n (2^{|A_i|} - 1)$, computing optimal values via enumeration can only be efficient for small games.

5.1 Reciprocity and Ultimatum Games

We considered instances of the *reciprocity* and *ultimatum* games from [6], which are shown in Figure 3. In each case, player 1 chooses between making a *fair* (*f*) or *greedy* (*g*) proposal, and player 2 decides to *reject* (*r*) or *accept* (*a*) it. The rectangular boxes show the corresponding utilities of players 1 and 2, with player 1’s utility being above that of player 2. We present the games, as in [6], in extensive form, with the players’ decisions taken sequentially, but will treat them as simultaneous moves in a single NFPG. Otherwise, beliefs would no longer be local since player 2’s utility would depend on an earlier decision by player 1.




Fig. 3: Reciprocity and ultimatum games in extensive form. The utilities for players 1 and 2 are given in the top and bottom rows of each leaf node, respectively.

In both games, player i attempts to maximise the expectation of a utility function that depends on the chosen actions (α) and belief (b) and has the form:

$$u_i(\alpha, b) = \lambda_i(\alpha) + \theta_i \cdot \kappa_{ij}(\alpha, b) \cdot \kappa_{ji}(\alpha, b)$$

where λ_i is a *material* utility function (these are the utilities shown in Figure 3), κ_{ij} reflects player i 's *kindness* to player j (expected material payoff, which ranges from negative to positive) and $\theta_i \in \mathbb{R}_{\geq 0}$ is player i 's *reciprocity sensitivity*. This type of game was originally studied in the context of *fairness equilibria* [31], in order to model and investigate scenarios in which agents are willing to sacrifice material utility to help or punish others depending on how they think they are being treated.

The concept of kindness was introduced in [31] as a way to measure this type of feeling, and is calculated as the difference between the utility that player i believes player j will receive (given player i 's choice) minus the average of the minimum and maximum utilities player i believes player j could get for i 's other choices. For instance, in the reciprocity game, if player 1 chooses *fair*, κ_{12} equals $5 - 1/2 \cdot (5 + (9 \cdot \bar{p}_r + 1 \cdot (1 - \bar{p}_r))) = 2 - 4 \cdot \bar{p}_r$. As before, we use p_r for the probability of player 2 choosing action r and \bar{p}_r for player 1's belief as to this value.

Reciprocating kindness is expressed by the matching of signs of κ_{ij} and κ_{ji} . Thus, if, by adopting a particular strategy, player i is perceived to be unkind to player j , κ_{ij} will be negative, which will in return motivate player j to be unkind to player i so that the product of κ_{ij} and κ_{ji} is positive. A similar logic applies to when players are perceived to be kind.

As explained earlier in Example 1, when writing the matrices for players' utility values, we can assume that $\bar{p}_r = p_r$ in equilibrium and just express them as functions of the probability p_r . For the reciprocity game we thus have:

$$\begin{aligned}
 \mathbf{Z}_1^{\text{reciprocity}} &= \begin{matrix} & r & a \\ f & \begin{pmatrix} 5 + \theta_1 \cdot (-4) \cdot (2 - 4 \cdot p_r) & 5 + \theta_1 \cdot (4) \cdot (2 - 4 \cdot p_r) \\ 1 + \theta_1 \cdot (-4) \cdot (4 \cdot p_r - 2) & 9 + \theta_1 \cdot (4) \cdot (4 \cdot p_r - 2) \end{pmatrix} \\ g & \end{matrix} \\
 \mathbf{Z}_2^{\text{reciprocity}} &= \begin{matrix} & r & a \\ f & \begin{pmatrix} 5 + \theta_2 \cdot (-4) \cdot (2 - 4 \cdot p_r) & 5 + \theta_2 \cdot (4) \cdot (2 - 4 \cdot p_r) \\ 9 + \theta_2 \cdot (-4) \cdot (4 \cdot p_r - 2) & 1 + \theta_2 \cdot (4) \cdot (4 \cdot p_r - 2) \end{pmatrix} \\ g & \end{matrix}
 \end{aligned}$$

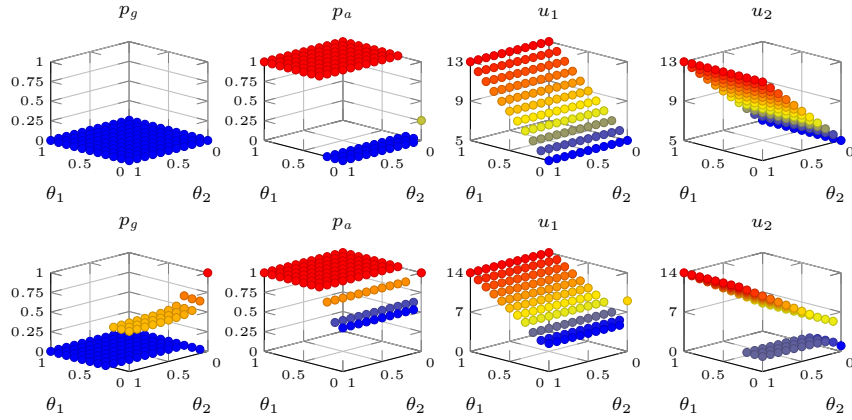


Fig. 4: Possible social welfare PEs for the reciprocity (top) and ultimatum (bottom) games for a range of reciprocity sensitivities θ_1, θ_2 . We show the strategy (probabilities p_g, p_a of choosing *greedy, accept*) and values (utilities u_1 and u_2).

and for the ultimatum game:

$$Z_1^{\text{ultimatum}} = f \begin{pmatrix} r & a \\ 5 + \theta_1 \cdot (-9/2) \cdot (2 + p_r/2) & 5 + \theta_1 \cdot (9/2) \cdot (2 + p_r/2) \\ 0 + \theta_1 \cdot (-9/2) \cdot (-2 - p_r/2) & 9 + \theta_1 \cdot (9/2) \cdot (-2 - p_r/2) \end{pmatrix}$$

$$Z_2^{\text{ultimatum}} = f \begin{pmatrix} r & a \\ 5 + \theta_2 \cdot (-9/2) \cdot (2 + p_r/2) & 5 + \theta_2 \cdot (9/2) \cdot (2 + p_r/2) \\ 0 + \theta_2 \cdot (-9/2) \cdot (-2 - p_r/2) & 1 + \theta_2 \cdot (9/2) \cdot (-2 - p_r/2) \end{pmatrix}$$

Figure 4 presents the strategies and utility values for SWNEs that we generated for the reciprocity and ultimatum games using different values of θ_1 and θ_2 . Although the games are very similar in structure (there is an equal amount of material utility that can be split by the two players in different ways), it is possible to see how the reciprocity sensitivity affects their behaviours and overall utilities. For instance, in the ultimatum game, when $\theta_1 = \theta_2 = 0$ the players are strictly concerned with their material utilities and (*greedy, accept*) is an acceptable SWNE as the sum of utilities is 10. It is possible to see though that, as θ_2 increases, player 1 is less likely to play *greedy* as the split becomes less fair for player 2, who could then retaliate by playing *reject*.

We can also notice that, when $\theta_1 = \theta_2 = 1$, the equilibrium for the reciprocity and ultimatum games is (*fair, accept*) which, despite leading to material utilities of (5, 5) in both games, accounts for different overall utilities for the players. In the former, the utility for each player is equal to $u_i(\text{fair, accept}) = 5 + 1 \cdot 4 \cdot (2 - 4 \cdot p_r) = 13$, and for the latter it corresponds to $u_i(\text{fair, accept}) = 5 + 1 \cdot 9/2 \cdot (2 + p_r/2) = 14$.




Fig. 5: An illustration for the pedestrian crossing scenario.

5.2 Traffic Games

We now report on a selection of case studies inspired by game-theoretic models of traffic and road user behaviour. We start with a simple one-shot game between a vehicle and a pedestrian in a road crossing scenario, and how their expectations can incentivise safe behaviour. Next we introduce a psychological variant of the Bayesian game presented in [29], which examined how cyclists would interact differently with autonomous and regular vehicles. Finally, we extend the road crossing scenario into a CSG and investigate the impact of combining information on past decisions with local expectations in a multi-stage, probabilistic model.

Pedestrian crossing. We consider a scenario where a pedestrian is deciding whether to cross a road near oncoming traffic, illustrated in Figure 5. We assume that the car has a right of way and can *reduce* (r) or *maintain* (m) its speed, while the pedestrian may choose to *cross* (c) or *wait* (w). A psychological game can be constructed by including incentives set to discourage behaviour based on what they expect the other will do. We assume the pedestrian would be (illegally) jaywalking if they decided to cross, and so give them a negative reward proportional to the probability p_c of that action being taken (multiplied by a constant μ), to model the pedestrian’s fear of being caught and incurring a penalty. This parameterisation results in the following utility matrices for the vehicle and the pedestrian:

$$Z_{vehicle} = \begin{matrix} & w & c \\ \begin{matrix} r \\ m \end{matrix} & \begin{pmatrix} 1-p_w & 1+p_c \\ 1+p_w & 1-p_c \end{pmatrix} \end{matrix} \quad Z_{pedestrian} = \begin{matrix} & w & c \\ \begin{matrix} r \\ m \end{matrix} & \begin{pmatrix} 1-p_r & 1+p_r-\mu\cdot p_c \\ 1+p_m & 1-p_m-\mu\cdot p_c \end{pmatrix} \end{matrix}$$

Figure 6 shows equilibria strategy profiles and utilities of this game for different values of μ (a more detailed version can be found in Figure 12, Appendix A). The colours for the points in the bottom and top plots for each value of μ




Fig. 6: Equilibria strategies and utilities for the vehicle (top) and pedestrian (bottom) in the pedestrian crossing scenario, for different values of μ .




Fig. 7: Original cyclist vs. vehicle game. In [29], the left and right games are played with probabilities p and $1 - p$, respectively. The utilities for the cyclist (1) and the vehicle (2) are given in the top and bottom rows of each leaf node.

serve to match the profiles and utilities of the pedestrian and the vehicle for different equilibria. While it is possible to see that there is always an equilibrium (displayed in red) in which the vehicle maintains its speed and the pedestrian waits, for smaller values of μ we also have an equilibrium strategy profile in which both agents make random choices. For example, in the equilibrium shown in blue for $\mu = 2$, the vehicle randomly chooses between reducing and maintaining its speed with probabilities $3/4$ and $1/4$ respectively, while the pedestrian crosses or not with probability $1/2$, potentially leading to unsafe behaviour. As the pedestrian’s uneasiness about crossing grows, i.e., as μ increases, the strategy profiles in which they cross with positive probability disappear, with the only remaining profile for $\mu = 5$ being the one in which they wait.

Cyclist vs. vehicle. Next, we model a cyclist and a vehicle, where the latter is either autonomous or driven by a human, at a road junction. A similar scenario was considered in [29], but modelled as a Bayesian game to investigate how increasing the share of autonomous vehicles affects the rate of collisions. Figure 7 shows the game in extensive form. The actions for the cyclist are *yield* (y), *walk* (w) and *cycle* (c), and the actions for the vehicle are *go* (g) and *stop* (s).

The utilities of the players reflect preferences over potential collisions, in accordance with traffic rules and an assumption on the part of the cyclist that an

α	u_{nature}	$u_{cyclist}$	$u_{vehicle}$
(a, y, g)	0	$5 \cdot p_a$	7
(a, y, s)	0	$3 \cdot p_a$	10
(a, w, g)	0	$-400 \cdot p_a$	-500
(a, w, s)	0	$15 \cdot p_a$	-15
(a, c, g)	0	$-500 \cdot p_a$	-300
(a, c, s)	0	$20 \cdot p_a$	15

α	u_{nature}	$u_{cyclist}$	$u_{vehicle}$
(h, y, g)	0	$8 \cdot p_h$	15
(h, y, s)	0	$6 \cdot p_h$	1
(h, w, g)	0	$-400 \cdot p_h$	-400
(h, w, s)	0	$15 \cdot p_h$	7
(h, c, g)	0	$-500 \cdot p_h$	-200
(h, c, s)	0	$20 \cdot p_h$	7

Table 1: Psychological cyclist vs. vehicle game in normal form. Nature chooses between *autonomous vehicle* (a) or *human driver* (h).




Fig. 8: Strategies and utilities for the cyclist vs. vehicle game.

autonomous vehicle (AV) would be programmed to be as *risk averse* as possible, while human drivers are often (unintentionally) distracted. So, for example, the penalty for an AV not stopping while the cyclist crosses the road is higher than for a human-driven vehicle (300 vs. 200). In the original game, it is assumed that a virtual *nature* player chooses the type of the vehicle according to a prior distribution. Here, we consider *nature* to be an active albeit *indifferent* player picking *autonomous* (a) or *human* (h), and we set the utilities for the cyclist to be the original utilities multiplied by their expectation of the vehicle being driven autonomously (p_a) or by a human (p_h), as detailed in Table 1.

Figure 8 shows results for this variant (a more detailed version is in Figure 11 of Appendix A). In addition to the equilibria computed in [29] (indicated in blue, orange³, magenta and brown), two new equilibria (indicated in green and red) are present for the psychological variant, in which *nature* randomly chooses between an autonomous vehicle and one with a human driver with probabilities $p_a = 14/17$ and $p_h = 3/17$. Considering the original model from [29] and combining the utilities of the two original games (Figure 7(a) and (b)) into one normal form game by multiplying the corresponding utilities by p or $1 - p$, we obtain the bimatrix game as follows:

$$Z_{cyclist} = \begin{matrix} & g & s \\ y & \begin{pmatrix} -3 \cdot p + 8 & -3 \cdot p + 6 \end{pmatrix} \\ w & \begin{pmatrix} -400 & 15 \end{pmatrix} \\ c & \begin{pmatrix} -500 & 20 \end{pmatrix} \end{matrix} \quad Z_{vehicle} = \begin{matrix} & g & s \\ y & \begin{pmatrix} -8 \cdot p + 15 & 9 \cdot p + 1 \end{pmatrix} \\ w & \begin{pmatrix} -100 \cdot p - 400 & 14 \cdot p + 7 \end{pmatrix} \\ c & \begin{pmatrix} -100 \cdot p - 200 & 8 \cdot p + 7 \end{pmatrix} \end{matrix}$$

³ The utility and strategy for the cyclist are the same as for the one in magenta.




Fig. 9: Transitions for actions (w, r) from the initial state of the CSG for the multi-stage pedestrian crossing example. States are of the form $\langle j, c_r, c_w \rangle$, for step count j and variables c_r, c_w .

We observe that, if the cyclist decides to *yield*, the vehicle would only play *go* with positive probability if $-8 \cdot p + 15 \geq 9 \cdot p + 1$, which gives $p \leq 14/17$. For any value of p above that threshold, the game has only one equilibrium in $(\textit{cycle}, \textit{stop})$, with utilities $(20, 8 \cdot p + 7)$. The analysis in [29] suggests that the number of collisions, i.e., the outcome where the cyclist chooses *cycle* and the vehicle chooses *go*, drops as the proportion of AVs increases. Indeed, if we only have AVs circulating, the only equilibrium (magenta) is $(\textit{cycle}, \textit{stop})$. However, by modelling this scenario as a three-player game with nature as an active player, we see in Figure 8 (and in more detail in Figure 11 of Appendix A), for a mix of AVs and human-driven vehicles, where AVs correspond to approximately 82% of the fleet, that the equilibrium strategy (green and red) for the cyclist is actually to *yield*. Furthermore, the vehicles can follow two different strategies: one in which they would *go* with probability 1 and another in which they *stop* with probability approximately 0.97 (shown in red and green, respectively). Finally, the model in [29] assumes that the cyclists can always differentiate between an AV and a human-driven vehicle, which is not realistic. In contrast, our model allows for the possibility of specifying psychological payoffs, meaning the actions taken by the cyclist can vary according to their beliefs about the type of the vehicle, paving the way for more sophisticated models of similar scenarios.

Multi-stage pedestrian crossing. Finally, we consider a probabilistic, multi-stage version of the earlier *pedestrian crossing* game, modelled as a CSG with psychological utilities. The state of the CSG has a variable j counting the number of times the one-shot pedestrian game has been repeated. We also add two discrete integer-valued variables $c_r, c_w \in \{0, 1, \dots, 10\}$ in order to carry information forward about the actions taken by both agents. The variables c_r and c_w are initialised to 0 and go up (or down) by 1 when the vehicle reduces (or maintains) its speed, and the pedestrian decides to wait (or cross), respectively.

To account for the fact that these observations can be imperfectly made by the players, c_r and c_w are updated probabilistically according to an *attention* coefficient $\gamma \in [0, 1]$, and so their values can also remain the same with probability $1 - \gamma$, as illustrated in Figure 9. The information on past decisions carried by the values of the two variables is then weighted with the local expectations

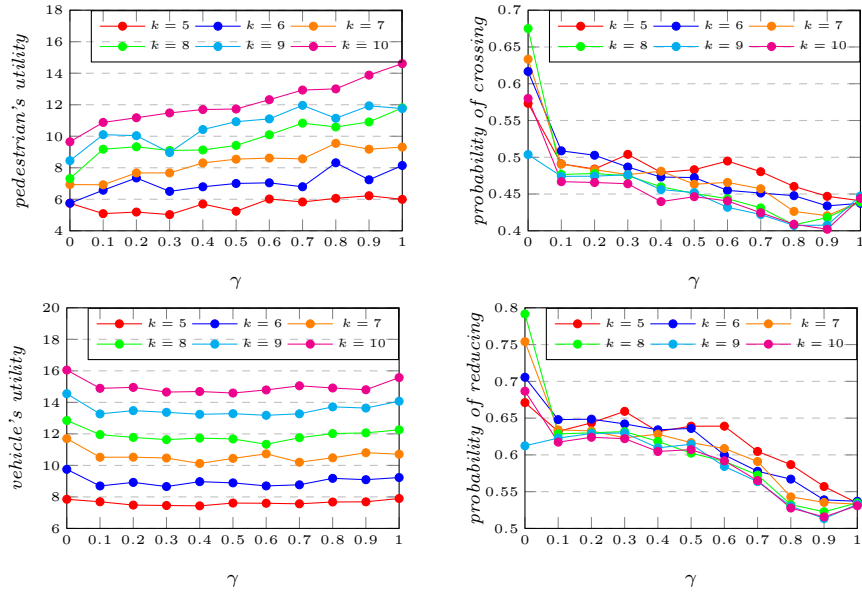


Fig. 10: Utility (left) and probability averages (right) over sampled equilibria for different path lengths in the multi-stage pedestrian crossing example.

at each state by modifying the reward matrices in following manner:

$$Z_{vehicle} = \begin{matrix} & \begin{matrix} w & c \end{matrix} \\ \begin{matrix} r \\ m \end{matrix} & \begin{pmatrix} 1-(1/2)\cdot(p_w+c_w/10) & (3/2)+(1/2)\cdot(p_c-c_w/10) \\ 1+(1/2)\cdot(p_w+c_w/10) & (1/2)-(1/2)\cdot(p_c-c_w/10) \end{pmatrix} \end{matrix}$$

$$Z_{pedestrian} = \begin{matrix} & \begin{matrix} w & c \end{matrix} \\ \begin{matrix} r \\ m \end{matrix} & \begin{pmatrix} 1-(1/2)\cdot(p_r+c_r/10) & 1+(1/2)\cdot(p_r+c_r/10)-\mu\cdot p_c \\ (3/2)+(1/2)\cdot(p_m-c_r/10) & (1/2)-(1/2)\cdot(p_m-c_r/10)-\mu\cdot p_c \end{pmatrix} \end{matrix}$$

Figure 10 shows expected utility values and crossing probability averages for different values of γ , paths of length k and $\mu = 1$. In a similar fashion to the one-shot example, we do not focus exclusively on the social welfare solution and there could be multiple equilibria for the NFGs built at each state. However, for extensive or multi-stage games, it is impractical to compute all possible equilibria unless the number of states is fairly small. This is particularly true when there is probabilistic branching, as the number of equilibria (and hence the number of NLPs to be solved) may grow very rapidly. For this reason, the probability values reported in Figure 10 (right) are averages over equilibria selected uniformly at random at each state. For each path length and γ , 10 experiments were run and the utility averages for the initial state are displayed in Figure 10 (left).

As the value of γ grows, it is possible to notice a trend in which the pedestrian's behaviour is safer, with the average probability of a crossing decreasing (top right). This is also reflected in higher utility for the pedestrian (top left),

$\gamma \in$	k	States	Transitions	Time(s)	k	States	Transitions	Time(s)
[0, 0]	5	6	21	0.76	8	9	33	1.33
(0, 1)		91	701	5.97		285	2,806	25.7
[1, 1]		91	256	4.82		285	897	23.1
[0, 0]	6	7	25	0.83	9	10	37	1.73
(0, 1)		140	1,198	9.12		385	3,981	39.6
[1, 1]		140	413	8.45		385	1,240	35.1
[0, 0]	7	8	29	1.06	10	11	41	2.09
(0, 1)		204	1,889	15.9		506	5,446	58.7
[1, 1]		204	624	14.5		506	1,661	51.9

Table 2: Statistics for the (CSG) multi-stage pedestrian crossing case study.

particularly when considering longer paths. For the vehicle, the utility averages (bottom left) remain fairly stable but we can also see that the likelihood of it reducing its speed (bottom right) decreases. This behaviour is desirable, considering it happens in coordination with decreasing probability values of a pedestrian crossing and the fact that the vehicle should have a right of way.

To provide an indication of model sizes that we are able to analyse, Table 2 shows how the number of CSG states and transitions, and the average computation time, vary for different values of k and intervals of γ (note that the number of transitions and states is the same for any value of $\gamma \in (0, 1)$, see Figure 9).

6 Conclusions

We have presented techniques that expand the scope of modelling and verification for game-theoretic probabilistic models. Starting with psychological normal form games, we proposed an NLP encoding that allows us to compute optimal equilibria for individual supports and, through support enumeration, an overall optimal equilibrium for a given NFPG. We then considered CSGs whose states can be expressed as NFPGs, and developed an algorithm to compute equilibria for such CSGs under some restrictions on the type of the players' beliefs. Finally, we reported on a prototype implementation and showcased novel automated analysis, made possible through our method, for a range of case studies.

Verification of psychological games is still a largely unexplored topic and there is ample room for expansion in theory, practice and applications to problems in computer science. Equilibria computation is a hard problem in general, and algorithms for psychological equilibria suffer from some of the same computational drawbacks as those for Nash or correlated equilibria, in addition to presenting new challenges of their own. The main current limitation is having to rely on enumeration for computing an optimal solution, which could be somewhat mitigated by parallelisation and filtering supports as a precomputation step. Future work includes investigating dynamic psychological games [5], which have the advantage of allowing belief updates but pose new modelling and computational challenges, and considering aspects of coordination and robustness via correlated [1] and trembling-hand [21] variants.

Acknowledgements. This project was funded by the ERC under the European Union’s Horizon 2020 research and innovation programme (FUN2MODEL, grant agreement No.834115).

References

1. Andrés, P.: From decision theory to game theory: Reasoning about decisions of others (2024), Manuscript in preparation, Cambridge University Press.
2. Aslanyan, Z., Nielson, F., Parker, D.: Quantitative verification and synthesis of attack-defence scenarios. In: Proc. CSF’16. pp. 105–119. IEEE (2016)
3. Attanasi, G., Nagel, R.: A survey of psychological games: Theoretical findings and experimental evidence. In: Innocenti, A., Sbriglia, P. (eds.) Games, Rationality and Behavior. Essays on Behavioral Game Theory and Experiments. pp. 204–232 (2008)
4. Battigalli, P., Corrao, R., Dufwenberg, M.: Incorporating belief-dependent motivation in games. *Journal of Economic Behavior & Organization* **167**, 185–218 (2019)
5. Battigalli, P., Dufwenberg, M.: Dynamic psychological games. *Journal of Economic Theory* **144**(1), 1–35 (2009)
6. Battigalli, P., Dufwenberg, M.: Belief-dependent motivations and psychological game theory. *Journal of Economic Literature* **60**(3), 833–82 (2022)
7. Bjørnskau, T.: The zebra crossing game – using game theory to explain a discrepancy between road user behaviour and traffic rules. *Safety Science* **92** (2015)
8. Brenguier, R.: PRALINE: A tool for computing Nash equilibria in concurrent games. In: Proc. CAV’13. LNCS, vol. 8044, pp. 890–895. Springer (2013)
9. Dickhaut, J., Kaplan, T.: A Program for Finding Nash Equilibria, pp. 148–166. Springer New York, New York, NY (1993)
10. Elvik, R.: A review of game-theoretic models of road user behaviour. *Accident Analysis & Prevention* **62**, 388–396 (2014)
11. Ezzati Amini, R., Abouelela, M., Dhamaniya, A., Friedrich, B., Antoniou, C.: A game-theoretic approach for modelling pedestrian-vehicle conflict resolutions in uncontrolled traffic environments. *Accident Analysis & Prevention* **203**, 107604 (2024)
12. Feng, L., Wiltsche, C., Humphrey, L., Topcu, U.: Synthesis of human-in-the-loop control protocols for autonomous systems. *IEEE Transactions on Automation Science and Engineering* **13**(2), 450–462 (2016)
13. Geanakoplos, J., Pearce, D., Stacchetti, E.: Psychological games and sequential rationality. *Games and Economic Behavior* **1**(1), 60–79 (1989)
14. Govindan, S., Wilson, R.: A global Newton method to compute Nash equilibria. *Journal of Economic Theory* **110**(1), 65–86 (2003)
15. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2021), [gurobi.com](https://www.gurobi.com)
16. Hensel, C., Junges, S., Katoen, J.P., Quatmann, T., Volk, M.: The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer* **24**(4), 589–610 (2022)
17. J. Gutierrez, J., Najib, M., Perelli, G., Wooldridge, M.: Eve: A tool for temporal equilibrium analysis. In: Proc. ATVA’18. LNCS, vol. 11138, pp. 551–557. Springer (2018), github.com/eve-mas/eve-parity
18. Junges, S., Jansen, N., Katoen, J.P., Topcu, U., Zhang, R., Hayhoe, M.: Model checking for safe navigation among humans. In: Proc. QEST’18. LNCS, vol. 11024, pp. 207–222. Springer (2018)

19. Karabag, M.O., Smith, S., Fridovich-Keil, D., Topcu, U.: Encouraging inferable behavior for autonomy: Repeated bimatrix stackelberg games with observations (2023)
20. Kemeny, J., Snell, J., Knapp, A.: *Denumerable Markov Chains*. Springer (1976)
21. Kolpin, V.: Equilibrium refinement in psychological games. *Games and Economic Behavior* **4**(2), 218–231 (1992)
22. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: *Proc. CAV’11*. LNCS, vol. 6806, pp. 585–591. Springer (2011)
23. Kwiatkowska, M., Norman, G., Parker, D., Santos, G.: PRISM-games 3.0: Stochastic game verification with concurrency, equilibria and time. In: *Proc. CAV’20*. LNCS, vol. 12225, pp. 475–487. Springer (2020)
24. Kwiatkowska, M., Norman, G., Parker, D., Santos, G.: Multi-player equilibria verification for concurrent stochastic games. In: *Proc. QEST’20*. LNCS, vol. 12289, pp. 74–95. Springer (2020)
25. Kwiatkowska, M., Norman, G., Parker, D., Santos, G.: Automatic verification of concurrent stochastic systems. *Formal Methods in System Design* **58**, 188–250 (2021)
26. Kwiatkowska, M., Norman, G., Parker, D., Santos, G.: Correlated equilibria and fairness in concurrent stochastic games. In: *Proc. TACAS’22*. LNCS, vol. 13244, pp. 60–78. Springer (2022)
27. Laan, G.V.D., Talman, A.J.J., Heyden, L.V.D.: Simplicial variable dimension algorithms for solving the nonlinear complementarity problem on a product of unit simplices using a general labelling. *Mathematics of Operations Research* **12**(3), 377–397 (1987)
28. Lemke, C., J. Howson, J.: Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* **12**(2), 413–423 (1964)
29. Michieli, U., Badia, L.: Game theoretic analysis of road user safety scenarios involving autonomous vehicles. In: *Proc. PIMRC’18*. pp. 1377–1381 (2018)
30. Porter, R., Nudelman, E., Shoham, Y.: Simple search methods for finding a Nash equilibrium. In: *Proc. AAAI’04*. pp. 664–669. AAAI Press (2004)
31. Rabin, M.: Incorporating fairness into game theory and economics. *The American economic review* pp. 1281–1302 (1993)
32. Sandholm, T., Gilpin, A., Conitzer, V.: Mixed-integer programming methods for finding Nash equilibria. In: *Proc. AAAI’05*. pp. 495–501. AAAI Press (2005)
33. Shapley, L.: Stochastic games. *Proc. National Academy of Sciences* **39**, 1095–1100 (1953)
34. Shirado, H., Kasahara, S., Christakis, N.A.: Emergence and collapse of reciprocity in semiautomatic driving coordination experiments with humans. *Proc. National Academy of Sciences* **120**(51) (2023)
35. Toumi, A., Gutierrez, J., Wooldridge, M.: A tool for the automated verification of Nash equilibria in concurrent games. In: *Proc. ICTAC’15*. LNCS, vol. 9399, pp. 583–594. Springer (2015)

A Appendix

Below, we include larger, more detailed figures for the *cyclist vs. vehicle* and *pedestrian crossing* case studies presented in Section 5.2.

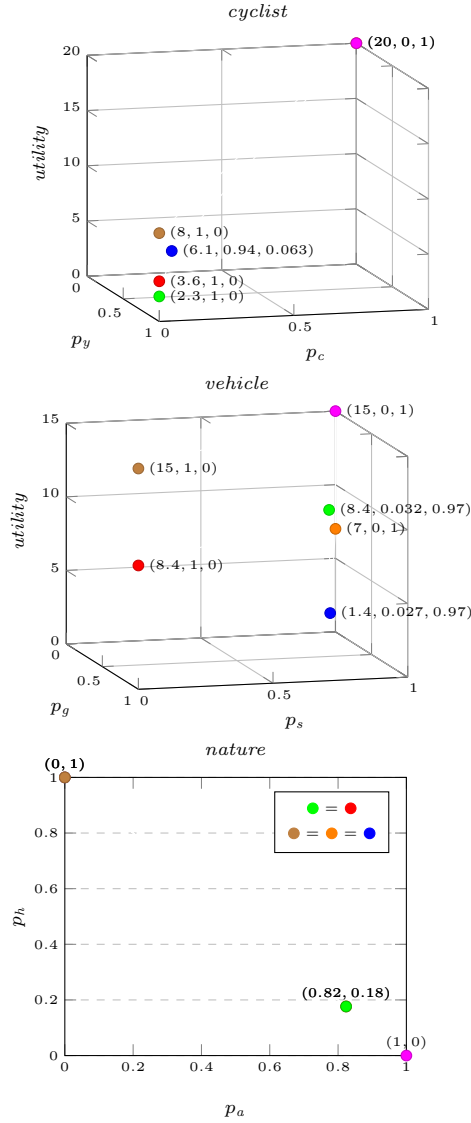


Fig. 11: Equilibria strategies and utilities for the cyclist vs. vehicle game. Coordinates for the *cyclist*, *vehicle* and *nature* players correspond to (u, p_y, p_c) , (u, p_g, p_s) and (p_a, p_h) , respectively (where u is the utility). For the cyclist, the utilities and strategies are the same in the orange and magenta equilibria.

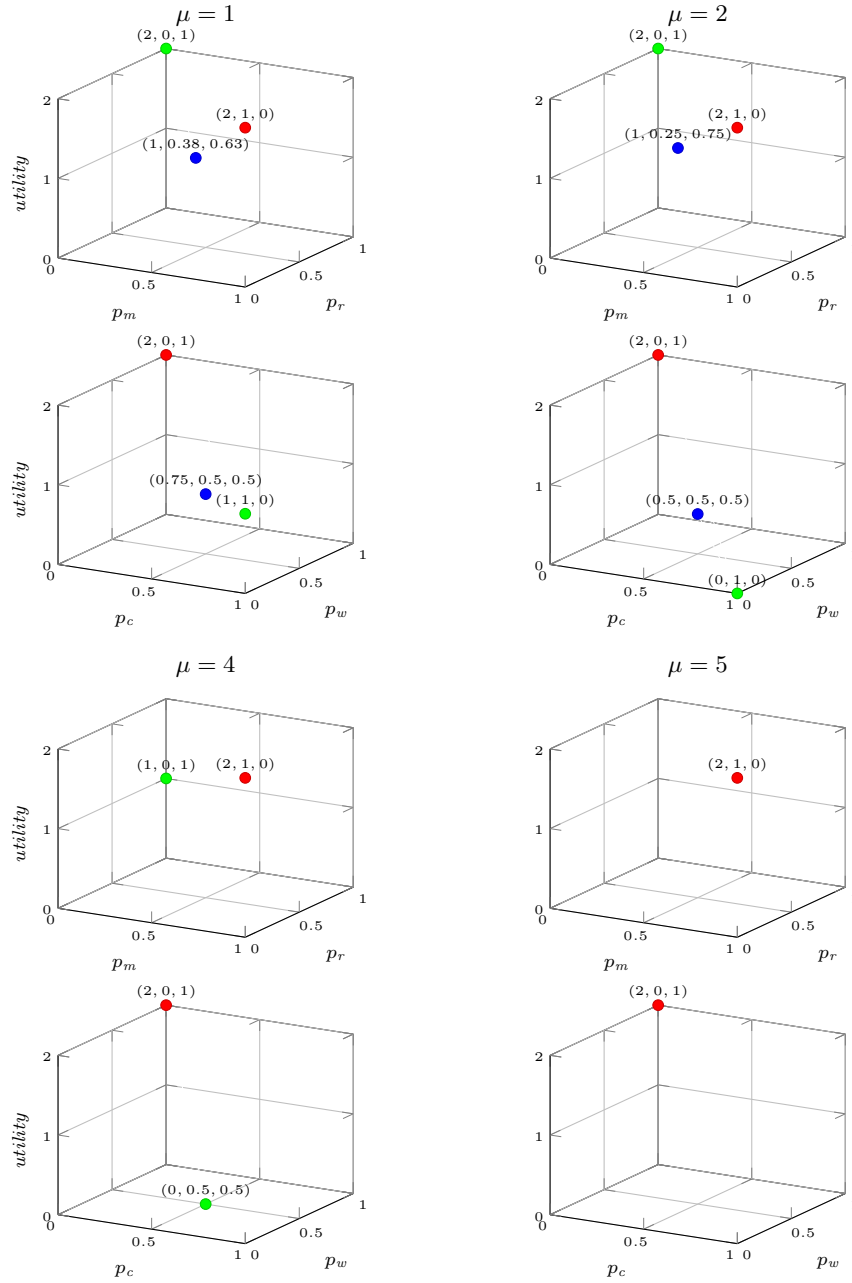


Fig. 12: Equilibria strategies and utilities for the pedestrian crossing game for different values of μ . The top plot for each value of μ corresponds to the pedestrian and the plot below to the vehicle. Point coordinates values are (u, p_m, p_r) for the vehicle and (u, p_c, p_w) for the pedestrian (where u is the utility).