CESifo WORKING PAPERS
9459
2021
December 2021

Procedurally Justifiable Strategies: Integrating Context Effects into Multistage Decision Making

Fynn Kemper, Philipp C. Wichardt

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University’s Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 9459*

Procedurally Justifiable Strategies: Integrating Context Effects into Multistage Decision Making

Abstract

This paper proposes a simple framework to model contextual influences on procedural decision making. In terms of utility, we differentiate between monetary payoffs and contextual psychological ones, e.g. deriving from the subjects' normative frame of reference. Monetary payoffs are treated as common knowledge while psychological payoffs are treated as partly unforeseeable. Regarding behaviour, we assume that players act optimal given their local perception of the game. As perceptions may be incorrect, we do not consider common equilibrium conditions but instead require strategies to be procedurally justifiable. As we will argue, various common inconsistencies considered in behavioural economics can be understood as procedurally justifiable behaviour. With the present framework, we add an abstract tool to the discussion which allows to consider also the behavioural implications of players foreseeing the corresponding behavioural effects - which is often not considered in the respective original models.

JEL-Codes: C700, D010, D910.

Keywords: behavioural inconsistencies, context effects, limited foresight, procedural decision making, utility.

Fynn Kemper
Department of Economics
University of Rostock
Ulmenstr. 69
Germany – 18057 Rostock
fynn.kemper@uni-rostock.de

Philipp C. Wichardt
Department of Economics
University of Rostock
Ulmenstr. 69
Germany – 18057 Rostock
philipp.wichardt@uni-rostock.de

This version: November 30, 2021
We are grateful to Leon Petrosyan and Denis Fedyanin for helpful comments and discussions. Kemper thanks the University of Rostock for financial support due to Landes-graduiertenförderung. The usual disclaimer applies.

1 Introduction

One of the fundamental assumptions in microeconomic theory is that every aspect that influences individual decision making can be expressed in units of utility. Put differently, following ones preferences means choosing what offers the higher utility. And, just as Bentham suggested when arguing for the general good as what creates “the greatest happiness of the greatest number” (cf. Bentham and Schofield, 2008), social welfare is commonly evaluated in terms of aggregate utility.[1] Moreover, to facilitate the analysis of economic decisions, utility is traditionally associated with an affine linear transformation of monetary incentives. If you have to decide for a pension plan as a consumer or a production plan as a firm, you will essentially be governed by the numbers.

However, as experimental economics has convincingly demonstrated over the last decades, focusing on monetary incentives alone is not enough once decisions are related to some form of prominent social context (see Camerer, 2003, for a review). As a consequence, a variety of behavioural models have been proposed to account for the various effects of fairness and reciprocity (e.g. Rabin, 1993; Charness and Rabin, 2002) other-regarding preferences (e.g. Levine, 1998; Fehr and Schmidt, 1999), regret (Loomes and Sugden, 1982), identity (Akerlof and Kranton, 2000) or a warm glow of giving (Andreoni, 1990). Yet, in this literature, it is common practice to focus on specific behavioural effects or very particular types of situations in which decisions are made and to adjust the modelling of utility accordingly.

The purpose of the present paper is to emphasise how many of these models eventually follow a similar abstract pattern. More specifically, we model utility as a combination of a monetary term and a non-monetary contextual component, referred to as psychological payoff (Bergh and Wichardt, 2018).[2] In doing so, we retain the standard assumption that monetary aspects of a decision are common knowledge. Yet, focusing on multistage interactions, we allow for the contextual component to be not perfectly foreseeable to decision makers. While the latter assumption seems plausible to make,[3] it will lead to non-existence of standard equilibria in strategic interactions due to expectations of one player (focusing on monetary incentives) being inconsistent with actual behaviour (due to "unforeseen" contextual effects).[4] Yet, different from standard models of strategic interactions, which treat unexpected actions that cannot be otherwise argued for as a sign of non-rational behaviour, we are more lenient with such inconsistencies. In the absence of perfect knowledge about contextual effects, we assume locally foreseen payoffs to be nothing but the best available

[1]Following Shackleton (1993) the likeliest source of the phrase as Bentham used it for the first time in 1776 is a translation of *Dei delitti e delle pene* published by Beccaria in 1768.

[2]Related ideas can also be found in Krupka and Weber (2013) who demonstrate the influence of identity effects on dictator game giving. For broader discussions of how the context may influence behaviour via identity see Akerlof and Kranton (2000) and Wichardt (2008).

[3]Can you say exactly ex ante how it will feel for you or even your colleague if you discontinue a joint project? Or if you refuse your child the desired ice cream? Especially relative to some fixed monetary value?

[4]Loewenstein et al. (2003) describe the projection bias which is the incorrect prediction of utilities in contexts different to the actual one. In a sense, we formalise the projection bias to model limited foresight of psychological payoffs. See Section 5 for further discussion.

guide regarding optimal behaviour; including the possibility for second thoughts later. Moreover, we assume that players are unaware of the fact that their (local) perception of the game need not be common knowledge.[5]

For the purposes of the argument, we use extensive games as the general basis and introduce an extension which we call *context game*. In doing so, we keep the common structure of histories but assume that the terminal payoffs, as perceived by the players, may vary as play moves on; a certain action may change the context in a way that induces evaluations of outcomes to change (the example discussed in the next section highlights this case). In that sense, our discussion is related to the literature on limited foresight in games (e.g. Jehiel, 2001; Jehiel and Lilico, 2010; Wichardt, 2010), which is concerned with the analysis of games in which players do not perfectly foresee the full structure of the game. In the present framework, however, limited foresight refers only to the effect the context has on preferences. All other game's details are assumed to be common knowledge. As we will see, if we assume perfect foresight and complete information, every context game can be shown to be solvable by backward induction; i.e. we retain a notion of subgame perfect equilibrium (the technical aspects of this approach are similar to Grossi and Turrini, 2012). With imperfect foresight, however, equilibria may fail to exist.

Regarding the non-existence of equilibria, we want to emphasise that we do not consider the effects of learning in our model. Depending on the details of the interaction, a properly defined learning procedure may well lead to perfect foresight and complete knowledge of contextual effects. Yet, similar to the literature on bounded rationality in general and the modelling of limited foresight in games in particular, we do not intend to model how players may get more rational.[6] Instead our focus is on understanding how games may evolve if players are subject to certain (arguably plausible) restrictions. Accordingly, while not able to rely on equilibria as mutually correct expectations, we will refer to the outcomes of play as *procedurally justifiable* if there is an ex post line of argument that rationalises - within the model - how the outcome came about. Moreover, we use the proposed framework to argue that many of the effects discussed in behavioural economics over the last decades, in fact, can be viewed as cases of such procedurally justifiable behaviour.

Before we move on to the details, however, we illustrate the key features of our framework by means of a simple technical example in Section 2. The abstract model, then, is defined in Section 3. As an illustration of our concept, two prominent examples are discussed in Section 4. In Section 5, we show how our framework relates to various existing models from behavioural economics and discuss how the multistage perspective in connection with the players' foresight may add to the understanding of the respective effects. Section 6 concludes.

[5] In a sense, the model entails a form of unawareness. We do not focus on the consequences of this aspect, though. See Heifetz et al. (2008), or Halpern and Rêgo (2014) for a more detailed discussion of unawareness.

[6] Personally, we have no doubt that people in general will tend to adjust behaviour towards better (expected) outcomes. Yet, we are sceptical regarding such processes really reaching equilibrium. Yet, we believe that a lot can be learned from attempts to "rationalise" what is observed.

2 Introductory Example

Consider the game depicted in Figure 1. There is only one player for this game, say Anne, who faces a two stage decision problem. At each stage, Anne can choose between two actions: A and B at stage one, and L and R at stage two, which is only reached after A. The game has three terminal nodes: z_1 after B, z_2 after AL, and z_3 after AR.

The general type of phenomena we aim to cover in the following are such for which preferences over outcomes change as the play of the game unfolds, i.e. we will allow the player's perception of utilities (preferences over outcomes) to be history-dependent. In the present setting, this means that Anne's preferences over the terminal histories z_2 and z_3 depend on whether she is still at h_0 or has already moved on to h_1; reflected in the two subgames with corresponding utility functions.

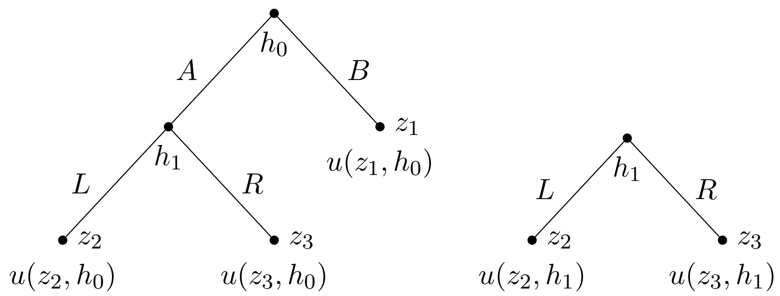

Figure 1: Game G_1: a two-stage decision problem for Anne. The left-hand game tree shows Anne's preferences at h_0, while the right-hand game tree reflects her preferences at h_1, i.e. once a decision at h_0 has been made.

To see why such cases, which are difficult to cover with standard game theoretic tools, are relevant in daily life decision making, think of Anne as considering to take her son, Erik, on her shopping tour (A); we will relate our argument to the more theoretical discussion in Section 3. Moreover, assume that Anne knows that Erik will ask for sweets once the till is reached but that she, by all means, does not want to give Erik any. However, while still at home, Anne is convinced to be able to withstand any requests at the shop (L) even in the face of other people looking, so that she prefers taking Erik to convincing him to stay at home (B). Once at the shop, however, Anne's forecast about her self-control turns out to be flawed, she feels bad about Erik making a scene and so ends up buying a huge package of sweets (R).

In terms of the formal decision problem presented in Figure 1, Anne prefers z_2 over z_1 and z_1 over z_3 at the initial history h_0. However, once the game has moved on, i.e. Anne and Erik have arrived at the shop, the local stimuli (people watching, social norms, self-control issues...) affect Anne more than expected so that now she prefers z_3 over z_2; see Figure 2 for illustration.[7]

[7]Note that, for the illustration, we have chosen the effect on preferences to be such that withstanding the

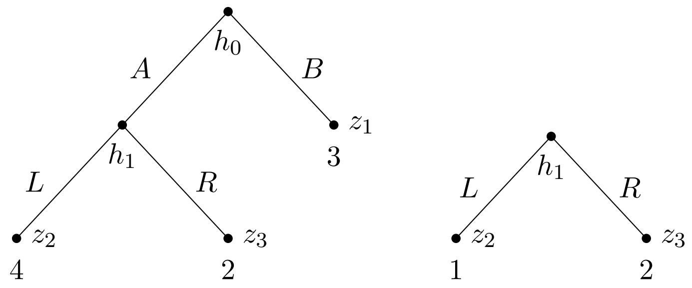

Figure 2: Game G_1: Anne's shopping experience.

What is crucial for the story as told above is that Anne is limited in her ability to correctly assess her future feelings - the social pressure etc. - when in the shop. For the illustrative purposes here, it is irrelevant whether she considers at all that her preferences may change or just underestimates the extent of the change. What is important is that Anne's preferences change as the game proceeds and that she would have decided differently at the outset – had she correctly foreseen the change.

Formally, again expressed in the notation of Figure 1, the general structure we see is that preferences at h_0 are given by

$$u(z_2, h_0) > u(z_1, h_0) > u(z_3, h_0),$$

while, due to context effects, i.e. circumstances that go beyond the strategic structure of the game, preferences at h_1 are

$$u(z_2, h_1) < u(z_3, h_1).$$

Thus, eventually Anne turns out at z_3 - the worst outcome in terms of her initial preferences. What is more, her doing so need not be seen as irrational (in the sense of not consistent with her preferences). While the value of monetary rewards may be easy to assess ex ante, psychological components connected with different outcomes often are not. We believe that the above example, albeit perhaps not of the utmost economic relevance, describes a common case in point. So, while the choice of a certain action may be inconsistent with initial plans, it need not be a sign of Anne being, for example, a crazy type (cf. Kreps et al., 1982).

In view of the various discussion of bounded rationality, we believe that psychological components of preferences create interesting cases where different levels of foresight are likely to exist and to be decisive for the eventual course of the game (for a discussion of foresight in games, see e.g. Jehiel, 2001; Jehiel and Lilico, 2010; Wichardt, 2010). A player with perfect foresight, for instance, may be able to consider the future influence of local contexts (histories) on preferences for the whole game tree, which would enable him to make

demands of Erik is much worse than expected implying that ex post leaving him at home would have been better. This choice, of course, is only one possibility. Other options are possible and will be discussed later.

a consistent plan of action at his first decision history by using backward induction.[8] By contrast, with limited foresight, a player is not fully aware of possible changes in preferences so that inconsistencies may arise; i.e. a player may make a plan which he will not necessarily follow in the future - as Anne in the above example.

In general, lack of foresight about ones own but also others' contextual payoffs may lead to various (seeming) inconsistencies and, hence, may render it difficult to find equilibria in strategic interactions in the conventional sense. Taking up the question how to model behaviour in such instances, we define the notion of a *procedurally justifiable strategy* below as one possible suggestion.

3 The Model

As a first step, we introduce what we will refer to as context games. The definition builds on extensive games with perfect information (see Osborne and Rubinstein, 1994, Chapter 6) and perfect recall but entails history-dependent payoffs. In a second step, we introduce a solution concept for context games with perfect foresight in analogy to subgame perfection in extensive games. Finally, the case of limited foresight is considered.

3.1 Context Games

Consider an extensive form game $G = (N, \mathcal{H}, P, \{u_i(z,h)\}_{i \in N, z \in Z, h \in \mathcal{H} \setminus Z})$, where N denotes the set of players, $\mathcal{H}$ the set of histories – with terminal histories collected in the set Z – and $P : \mathcal{H} \setminus Z \to N$ assigns to each $h \in \mathcal{H} \setminus Z$ the player who is about to move. Player $P(h)$, then, chooses a probability distribution $\sigma_i(h)$, $i = P(h)$, over all possible continuations, i.e. histories following h – also referred to as actions a from the set of available action at history h, denoted by $A(h)$. Moreover, as usual, a player i's (behaviour) strategy σ_i is the collection of all his choices in the course of the game. All possible strategies of player i, $i \in N$, are collected in the set ΔS_i, where ΔS_i refers to mixed and S_i refers to pure (behaviour) strategies. Finally, for all $i \in N$, player i's preferences over terminal nodes $z \in Z$ - as evaluated at history $h \in \mathcal{H}$ - are described by the utility function $u_i : Z \times \mathcal{H} \to \mathbb{R}$.

Note that in the present setting utility functions are history-dependent, i.e. preferences over outcomes may change during the game.[9] More specifically, we think of utility as consisting of two parts: A monetary component $m_i : Z \to \mathbb{R}$, referred to as *monetary payoff*, which depends only on the terminal history, and a context-dependent part $\pi_i : Z \times \mathcal{H} \to \mathbb{R}$,

[8] Here, "consistent" means that the plan will be actually followed as the game proceeds. This understanding goes back to Strotz (1955).

[9] Here, we do not view utility in a consequentialist perspective as it is typically done in economics (cf. Loewenstein et al., 2001). Psychological payoffs are meant to integrate also context effects, i. e. anticipatory emotions, that are only present at the time of decision making - like fear induced by a risky choice. In this sense, utility in our model is somewhat similar to the notion of decision utility proposed by Kahneman (1994).

referred to as *psychological payoff*, which is thought of as covering non-monetary aspects of the decision making process. Thus, the context, i.e. the state of the game, becomes a parameter of the utility function. Player i's expected utility at non-terminal history $h \in \mathcal{H} \backslash Z$ under the strategy profile $\sigma = (\sigma_i, \sigma_{-i}) \in \Delta S := \times_{i \in N} \Delta S_i$, then, is given by

$$EU_i(\sigma, h) = \sum_{z \in Z} O(\sigma)(z) u_i(z, h),$$

where $O(\sigma)$ refers to the probability distribution over the set of terminal histories Z as induced by the strategy profile σ, i.e. $O(\sigma)(z)$ is the probability of the terminal history z under the strategy profile σ; $O(\sigma)$ is referred to as the outcome of σ. For the time being, we leave open the question about what exactly player i at h knows (or assumes) about σ_{-i}. We will later return to this issue and discuss how answering it depends on the players' perception of the game regarding future non-monetary components of utility. Independent of the exact details of the answer, we will assume that player i at node h always conceives of themselves and others as following a plan of action that is optimal for them given what player i knows about their own and other players' payoffs at node h.[10]

For the sake of argument, we consider an additively separable utility function, a common choice in behavioural economic models (e.g. O'Donoghue and Rabin, 1999, Levine, 1998, Fehr and Schmidt, 1999), i.e.:

$$u_i(z, h) = m_i(z) + \pi_i(z, h).$$

A *context game*, then, is a standard extensive form game in which the locally perceived values of terminal utility depend on the state of the game.

3.2 Foresight

In the following, we focus on two cases regarding utility: (a) one in which all players at all instances are perfectly informed about the game's payoff structure (perfect foresight) and (b) one in which players at any point of the game, i.e. any $h \in \mathcal{H} \backslash Z$, have an imperfect understanding of their own and others' subsequent psychological payoffs (limited foresight).

In order to model these cases, we say that player i, $i \in N$, at history $h \in \mathcal{H} \backslash Z$ perceives the psychological payoff of player j, $j \in N$, at node $z \in Z$ as evaluated by player j at history h', i.e. $\pi_j(z, h')$, as $\pi_j^i(z, h' | h, \delta_i)$; $\delta_i \in [0, 1]$ is referred to as the foresight parameter of player i. Moreover, we say that player $i \in N$ has *perfect foresight* if $\delta_i = 1$, i.e. $\pi_j^i(z, h' | h, \delta_i = 1) = \pi_j(z, h')$ and *limited foresight* if $\delta_i < 1$. The vector of the players' foresight parameters is denoted by $\delta = (\delta_i)_{i \in N}$.

[10] Of course, other assumptions - such as players assuming others to follow certain biases or repeating "earlier" mistakes - are possible. Eventually, what is best is an empirical question which will not be addressed in the present paper. The assumption made in the text seems most plausible to us. Yet, it is made only to be able to demonstrate the dynamics of the framework without any claim at (closeness to) truth.

A *context game with perfect foresight* corresponds to the case, where $\delta_i = 1$ for all $i \in N$, respectively $\delta = \mathbf{1}$. Otherwise, we refer to the game as a *context game with limited foresight.* By assumption, for every player i, the perception bias $|\pi_j^i(z, h'|h, \delta_i) - \pi_j(z, h')|$ decreases in δ_i. Hence, a higher level of foresight corresponds to perceptions that are closer to 'real' payoffs.

In order to keep the argument technically tractable, we assume that a player's perception of his own future payoffs is a weighted average of his actual payoff at that later history and his payoff dependent at his actual history. This can be seen as an application of the projection bias (Loewenstein et al., 2003), which is discussed in Section 5 in more detail. Moreover, we assume that the bias increases with the distance between the actual history h and the later one h'. Let $d(h, h')$ denote the number of actions it takes to reach history h' when starting at history h. Then, we define the value player i assigns at history h to his own psychological payoff from terminal history z as experienced at history h',

$$\pi_i^i(z, h'|h, \delta_i) = \delta_i^{d(h,h')}\, \pi_i(z, h') + (1 - \delta_i^{d(h,h')})\, \pi_i(z, h).$$

Note, that this function only refers to a player's perceptions about his own psychological payoffs. The modelling of a players' perceptions about another player's psychological payoffs is discussed in Remark 1 further below.

Assuming monetary payoffs to be perceived correctly, we obtain the following expression for player i's, $i \in N$, overall perception of utilities at history $h \in \mathcal{H} \setminus Z$.

$$u_j^i(z, h'|h, \delta_i) = m_j(z) + \pi_j^i(z, h'|h, \delta_i)$$

Put differently, $u_j^i(z, h'|h, \delta_i)$ represents how player i at history h assesses the utility player j at history h' assigns to terminal history z. Note, that we allow for $i = j$, i.e. a player may have a wrong perception of his own psychological payoffs at later histories (as in the Example in Section 2); however, for all $i \in N$ we have $\pi_i^i(z, h|h, \delta_i) = \pi_i(z, h)$.

The solution concept for context games with limited foresight is based on subgame perfection (Selten, 1965). Following common standards, we define a subgame of a context game G starting at history $h \in \mathcal{H} \setminus Z$ as $G(h) = (N, \mathcal{H}|_h, P|_h, \{u_i|_h(z, h')\}_{i \in N, z \in Z|_h, h' \in (\mathcal{H} \setminus Z)|_h})$, where $\mathcal{H}|_h$ consists of all h' for which $(h, h') \in \mathcal{H}$. $Z|_h$ includes all terminal histories of $\mathcal{H}|_h$ and $(\mathcal{H} \setminus Z)|_h$ describes all non-terminal histories of $\mathcal{H}|_h$. The player function $P|_h$ is given by $P|_h(h') = P(h, h')$ for all $h' \in \mathcal{H}|_h$. A strategy of player $i \in N$ in the subgame $G(h)$ is denoted by $\sigma_i|_h$; it is derived by the strategies of the original game, such that $\sigma_i|_h(h') = \sigma_i(h, h')$ with $h' \in \mathcal{H}|_h$. The set of all possible strategy profiles $\sigma|_h = (\sigma_i|_h, \sigma_{-i}|_h)$ is denoted by $\Sigma|_h = \times_{i \in N} \Sigma_i|_h$. The players' utility functions $u_i|_h$, $i \in N$, are defined by $u_i|_h(z, h') = u_i(z, (h, h'))$, for all $h \in \mathcal{H}$.

As perceptions may vary, we use the notion of a *perceived* subgame of a context game

$G_\delta(h) = (N, \mathcal{H}|_h, P|_h, \{u_j^{P(h)}|_h(z, h'|h, \delta_{P(h)})\}_{j\in N, z\in Z|_h, h'\in(\mathcal{H}\backslash Z)|_h})$. It is defined such that $G_\delta(h)$ has the same structure as the subgame $G(h)$ except that utilities are replaced by the corresponding perceptions of the respective player $P(h)$ which are the result of hist (limited) foresight. Note that in the case of perfect foresight, i.e. $\delta = \mathbf{1}$, the perceived subgame of a context game is the original subgame, $G_\delta(h) = G(h)$. Moreover, we assume that every player is unaware of his bias in perceptions but, whenever called upon to move, believes that his perceptions reflect true utilities and are common knowledge.

Intuitively, perceived payoffs may be understood as the players' best guess about actual payoffs at succeeding histories. As the game moves, on these perceptions may turn out to be wrong, though. Yet, without learning, local perceptions what guides behaviour. From a technical perspective, this poses no further problem as a perceived subgame of a context game is itself a context game and can be analysed as such.

3.2.1 Perfect Foresight

Next, we present a solution concept for context games with perfect foresight, which means that players have perfect knowledge of all the game's details. In analogy to Grossi and Turrini (2012), we use a notion of subgame perfection. In our framework, each player considers the history-dependent payoff structure of the subgame that starts at his specific decision history. Then, each player at each of his decision histories maximises his payoff depending on his perception of the game at this specific history. The result is described by the definition below, which applies the one deviation property.

Definition 1. *A strategy profile σ^* in a context game G with perfect foresight is called 'subgame perfect' if for every player $i \in N$ and every history $h \in \mathcal{H}\backslash Z$ with $P(h) = i$:*

$$EU_i(\sigma_i^*|_h, \sigma_{-i}^*|_h, h) \geq EU_i(\sigma_i'|_h, \sigma_{-i}^*|_h, h) \qquad \text{for all } \sigma_i'|_h \in \Sigma_i|_h,$$

where $\sigma_i'|_h$ differs from $\sigma_i^*|_h$ only in the action it prescribes at history h.

As the definition coincides with standard subgame perfection in all relevant aspects, it is immediate that such a strategy profile always exists.

Proposition 1. *Every context game G with perfect foresight (with generic payoffs) has a (unique) subgame perfect strategy profile.*

Proof. Analogue to the proof of Kuhn's Theorem (Kuhn, 1953). □

It is worth noting that the subgame perfect strategy profile might not be unique for non-generic context games. However, as payoffs are understood as a combination of a monetary and a psychological component, non-generic cases seem highly unlikely in view of applications. In fact, while in real life situations monetary payoffs may indeed be identical, we consider it to be highly implausible that psychological payoffs are. Accordingly, the resulting subgame perfect strategy profile is unique for almost all relevant context games with perfect foresight.

A further point to note is that subgame perfect strategy profiles need not lead to ex ante optimal outcomes - as already exemplified in the introductory example. The reason being that players are aware of later changes in their assessments and incorporate these in their behaviour (in terms of the example: expecting she will loose her self-control, Anne will not take Erik to the shop, despite it feeling optimal at home).

Proposition 2. *A subgame perfect strategy of a context game with perfect foresight need not implement the highest payoff as perceived at the initial history.*

Proof. See introductory Example in Section 2. □

3.2.2 Limited Foresight

Next, we move on to the case of restricted foresight regarding context effects. With respect to applications, this in fact seems to be the most plausible case. In particular, while monetary rewards are often comparably clear for different outcomes – so that even assuming them to be common knowledge in cases with more than one player is arguably innocuous – psychological payoffs are likely to be far more fickle and difficult to foresee.

Of course, context games with limited foresight regarding psychological payoff components give rise to various technical difficulties. In particular, we need to cover the inconsistencies that may arise between the players' perceptions and eventual 'real' payoffs. For example, in general, two players i and j will not share the same perception regarding player k's payoff at history h', $u_k^i(z, h'|h, \delta_i) \neq u_k^j(z, h'|h, \delta_j)$. This holds even if their foresight levels are the same, i.e. $\delta_i = \delta_j$.

In order to be able to analyse such situations, we consider a solution concept that treats the individual players' perceptions as all that matters for the decision making of a particular player at a particular history. The concept is based on the modified notion of subgame perfection as presented in Definition 1, the main difference being that due to the possibility of mismatching subjective perceptions of the game, we cannot expect "equilibrium" in the sense of correct expectations of behaviour (neither for oneself nor for others).

Remark 1. *Note that a standard correct expectations approach would require players to make (correct) predictions about their own perceptions of future contextual effects as well as the perceptions of and about other players' perceptions etc. - in addition to correct assessments of other players' preferences in general. To us, this seems highly implausible to achieve. In fact, we see the literature on framing effects as - among other things - demonstrating that even the scientific experts are not perfect in foreseeing how a certain context will affect the respective decision makers (e.g.Cookson, 2000; Bergh and Wichardt, 2018). Moreover, there* *is a whole literature in psychology on mentalisation / theory of mind, which addresses how we conceive of others in a way that apparently differs from "correct expectations" (see Apperly, 2012, and Katznelson, 2014, for reviews and further references).*[11]

For the present argument, we try to capture the corresponding more self-centred view and drop standard equilibrium conditions. Instead, at each non-terminal history, we consider only perceptions of the moving player (which are needed to assess optimality of own and others present and future behaviour at that history). Moreover, we consider only perceptions regarding succeeding histories; i.e. for some player $P(h) = i$ *only those perceptions* $\pi_j^i(z, h'|h, \delta_i)$ *are needed for which* h' *is a successor of* h*. The reason for this is that our solution concept is based on* perceived *subgames, which only entail the perception of the specific player who is in charge of the first action in that subgame. These perceptions alone build a subgame which is treated like a context game with perfect foresight. All subgame perfect strategy profiles of all perceived games together, then, are the basis of the procedurally justifiable strategy profile.*

To begin with, recall that $G_\delta(h)$ (in general) differs from $G(h)$, $h \in \mathcal{H}\backslash Z$, due to the players' limited foresight. In particular, depending on their foresight level, players will have more ($\delta_i \to 0$) or less ($\delta_i \to 1$) blurred perceptions of the non-monetary consequences involved in the strategic interaction. Thus, $G_\delta(h)$ describes how player $P(h)$, who is subject to limited foresight $\delta_{P(h)}$, perceives the subgame starting at history h. Different from the case of perfect foresight, however, players do not (correctly) foresee how they will assess non-monetary consequences of behaviour once the game has moved past h.

Accordingly, regarding strategy choices, player $P(h)$ uses his (subjective) information about $G_\delta(h)$ as available at h to determine his optimal behaviour at h, i.e. $\sigma_{P(h)}(h)$. In doing so, we assume that he solves $G_\delta(h)$ by backward induction to derive a (genericly unique) subgame perfect strategy profile for $G_\delta(h)$, $\tilde{\sigma}^*(h)$ (in non-generic cases, player can choose any subgame perfect profile). Thus, we require that player $P(h)$ assumes a behaviour of other players which, given his perspective on G at node h, would constitute a subgame perfect equilibrium for a game $\tilde{G}$ with payoffs as given by $G_\delta(h)$ and perfect foresight if combined with his $\sigma_{P(h)}(h)$.

Once optimal decisions are determined for all histories $h \in \mathcal{H}\backslash Z$ and respective players $P(h)$, we say that the resulting probability distribution over terminal nodes is *procedurally justifiable* if at all histories players act optimally given their local perception of the game. Note, however, that what is perceived as optimal at a certain node by a given player may change as the game moves on (as his perception may change).

[11] Van Boven and Loewenstein (2005) pointed out that perspective taking depends on how similar to oneself the other individual is perceived to be. If the other person is assumed to be similar to oneself, one applies some kind of egocentrical perspective taking - first one estimates how oneself would feel in the situation of the other person, second this estimation is changed in order to account for dissimilarities between oneself and the other person. When considering people who are perceived as different to oneself, perspective taking is based on stereotypes instead.

Definition 2. *Given a context game G with limited foresight and foresight levels $\delta = (\delta_i)_{i \in N}$. For any history $h \in \mathcal{H} \backslash Z$, let $\tilde{G}(h)$ be the game derived from $G_\delta(h)$ by assuming payoffs to actually are as perceived at h by $P(h)$ and perfect foresight. A strategy profile σ^* in G is called* procedurally justifiable *if for every player $i \in N$ and every history $h \in \mathcal{H} \backslash Z$, we have:*

$$\sigma^*_{P(h)}(h) = \tilde{\sigma}^*_{i,h}(h), \text{ where } \tilde{\sigma}_{i,h} \text{ is part of some subgame perfect strategy profile } \tilde{\sigma}^*(h) \text{ of } \tilde{G}(h).$$

The corresponding outcome $O(\sigma^*) \in Z$ is called a procedurally justifiable outcome*.*

Remark 2. *Note that subgame perfect strategies and procedurally justifiable strategies are identical for generic context games if δ_i is sufficiently close to 1. This is an immediate consequence of the fact that, by assumption, perception biases decrease continuously in δ. In particular, for every δ there exists an left-open interval with upper limit 1, i.e. the perfect foresight case, such that for all δ from this interval the players' choices under limited foresight are the same as in the perfect foresight case. Thus, if foresight is limited only slightly, players end up with the same strategy profile as if foresight was not limited at all.*

Existence and uniqueness (for generic games) of the procedurally justifiable outcome follow immediately from the construction as the following proposition shows.

Proposition 3. *Every context game G with limited foresight (and generic payoffs) has a (unique) procedurally justifiable strategy profile.*

Proof. By construction. By Proposition 1 every context game with perfect foresight has a (unique) subgame perfect strategy profile. Furthermore, for every history $h \in \mathcal{H} \backslash Z$ the game $\tilde{G}(h)$ induced by $G_\delta(h)$ is a context game with perfect foresight (by definition). Accordingly, Proposition 1 applies and the claim of Proposition 3 follows. □

Recall that a procedurally justifiable strategy profile need not satisfy common equilibrium requirements as false expectations about own and other players' payoffs may lead to false beliefs about own and other players' (later) actions. While it is common to consider equilibrium behaviour in economics, we believe that the desire to study psychological effects in a more technical framework requires compromises.

Example

In order to exemplify our solution concept, we return to the introductory example from Section 2. We denote the described context game from Section 2 by $\hat{G}$. The corresponding perceived subgame of Anne at h_0, $\hat{G}_\delta(h_0)$, is depicted in Figure 3. Here, we assume that Anne's foresight is limited, i.e. $\delta < 1$, so that $\pi(z, h'|h, \delta) = \delta^{d(h,h')}\, \pi(z, h') + (1 - \delta^{d(h,h')})\, \pi(z, h)$ as introduced in Section 3.1 in order to derive the necessary perceptions of future psychological payoffs. Note that for the sake of argument, Anne's player index is suppressed in sequel.

The first point to note is that if δ is small enough the preference relation $u(z_2, h_0) > u(z_3, h_0)$ carries over to $u(z_2, h_1|h_0, \delta) > u(z_3, h_1|h_0, \delta)$. Accordingly, the subgame perfect strategy of $\tilde{G}_\delta(h_0)$ is (AL). Moreover, the perceived subgame in h_1 equals the right-hand side of Figure 1 as there is only the initial history h_1 and corresponding payoffs are perceived correctly. Hence, the unique subgame perfect strategy for this subgame is (R). It follows that (AR) is the only procedurally justifiable strategy of $\hat{G}$ for small δ, so that in the end Anne takes Erik on her shopping tour and buys him the sweets he asks for.

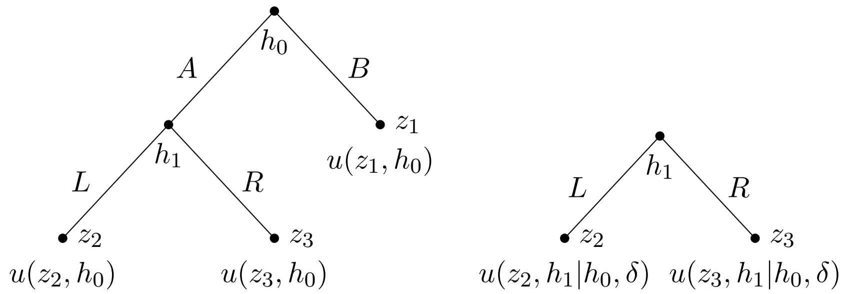

Figure 3: $\hat{G}_\delta(h_0)$: the game from the introductory example as it is perceived from h_0 with limited foresight.

While the above discussion suggests that more foresight (higher δ_i) is better as it is closer to rationality, the following proposition shows that this need not be the case; i.e. there are cases where limited foresight is advantageous (the effect of limited foresight being similar to that of a commitment device).

Proposition 4. *There exist context games with limited foresight for which the procedurally justifiable outcome yields a higher payoff under limited foresight than under perfect foresight.*

Proof. As illustrated by the following example using the context game $\tilde{G}$ from Figure 4.[12]

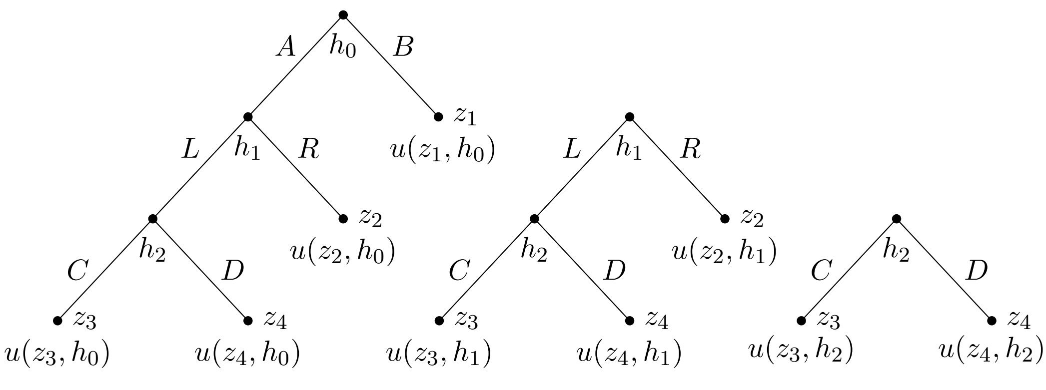

Figure 4: Context Game $\tilde{G}$: A three-stage context game with one player.

[12]This type of result is also presented by O'Donoghue and Rabin (1999). We will come back to its original version in Section 5.

Let the player's payoffs dependent on the initial history be such that $u(z_3, h_0) > u(z_4, h_0) > u(z_1, h_0) > u(z_2, h_0)$. Moreover, assume that he has preferences represented by $u(z_3, h_1) > u(z_2, h_1) > u(z_4, h_1)$ at history h_1 and $u(z_4, h_2) > u(z_3, h_2)$ at history h_2.

An agent with $\delta = 1$ has correct perceptions about his future payoffs, i.e.$u(z, h'|h, \delta = 1) = u(z, h')$. At h_0 he knows that D would be played if h_2 was reached as $u(z_4, h_2) > u(z_3, h_2)$, so that the outcome would be z_4. He also knows, that he will have this knowledge at h_1 when he technically has to decide between the outcomes z_2 (if R is played) and z_4 (if L is played). Moreover, his preferences at h_1 are $u(z_2, h_1) > u(z_4, h_1)$, so that R will be chosen, if h_1 is reached. By a similar argument, he chooses B at the beginning of the game as $u(z_1, h_0) > u(z_2, h_0)$. Hence, the procedurally justifiable strategy is (BRD) which leaves him with the outcome z_1.

In case of $\delta = 0$ and with the foresight function from Section 3.1 an agent's perceptions are given by $u(z, h'|h, \delta = 0) = u(z, h)$. Thus, the agent does not think of his preferences as history-dependent. According to his preferences at h_0, he expects to play C at h_2 and L at h_1, which would leave him with his best possible outcome with regard to his preferences at h_0. Therefore, he chooses A. Then, at h_1 his preferences have changed. As he does not consider any further preference change in case of reaching h_2, he expects that C will be played and, therefore, chooses L. However, at h_2 his preferences change again and he chooses D. Hence, the procedurally justifiable strategy is (ALD) for $\delta = 0$, which leads to outcome z_4. Thus, despite his wrong understanding of the game's dynamics, the player reaches a better outcome with regard to his preferences at h_0 than a player with perfect foresight would have done. □

4 Examples

The present section discusses two prominent empirical examples in order to illustrate how the concept of procedurally justifiable strategies may help to identify crucial motivations for the players.

4.1 The Marshmallow Experiment

As a first example, consider the famous marshmallow experiment by Mischel et al. (1989). In this experiment, various children were given the choice between an instantaneous treat or to wait some time and receive more treats. Some children chose the instantaneous gratification; others were able to resist - mostly by distracting themselves from the treats. Later, it was shown that these differences in behaviour are correlated with cognitive and academic competence in adolescence (Shoda et al., 1990).[13]

[13]Watts et al. (2018) found much weaker results with a more diverse sample of children and more sophisticated statistical models, though.

Consider the simplified marshmallow experiment depicted in Figure 5.[14] In this one player two stage game, the player first has to determine whether to look (at the reward) or to turn away. As turning away will need some self-control, we assume a small cost $\varepsilon > 0$ for it as well as a discount factor of $\xi \in (\frac{1}{2}, 1)$ for future rewards. Moreover, we assume that actually looking at the reward increases its perceived value from 1 to $1 + E > 2\xi$, i.e. there is a contextual benefit associated with grabbing the reward when focusing on it.

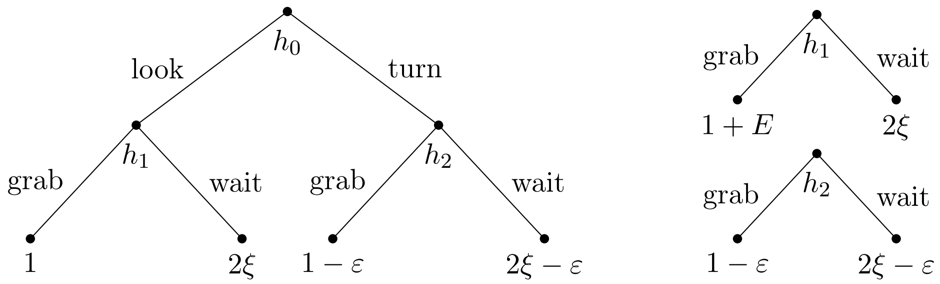

Figure 5: The simplified marshmallow experiment. Game trees on the right showing subgames upon being reached (E being reflection the temptation upon seeing the reward).

First, assume that the player has perfect foresight. Thus, already at h_0 he foresees that he would grab the one marshmallow at h_1, while after *turn* he would prefer to *wait* at h_2. Accordingly, *turn* and *wait* is what he will decide to do at h_0 and will follow at h_2. Hence, the subgame perfect strategy in the case of perfect foresight is (*turn grab wait*).

In oder to illustrate the case with limited foresight, we assume $\delta = 0$, the player has no idea about future changes in his payoffs. Accordingly, *look* and *turn* differ only in the extra cost of ε when evaluated at h_0. In this case, the unique procedurally justifiable strategy is (*look grab wait*). Moreover, for intermediate levels of foresight, what is decisive is whether the player is sufficiently aware of the change in his evaluations so that - based on his evaluations at h_0 - he will prefer to wait (as at h_0 two rewards are judged better than one).

Note that we do not make any assumptions about whether one or two rewards are actually better for the player as considered by him. The only thing the argument above says is that a player who (himself!) judges two rewards better than one at h_0 and who foresees that he will not stick to this at h_1 will act in a way at h_0 that guarantees him what he considers best at h_0. We emphasise this point as for the marshmallow experiment it is generally considered to be good to wait. Admittedly, adults typically think that waiting a little in order to get two desired items is worth the while. Whether the children involved in the decision actually think so, too, or not, we don't know, though. Hence, we make no such assumption.[15]

[14] The actual situation was roughly as follows: A child sits in a room with an experimenter and gets informed about the experiment's details. The child gets to see the possible rewards in front of them. The possibilities are either one marshmallow instantly or two marshmallows later if the child waits some time. The experiment begins after the child understood all information.

[15] The problem of which perspective is the appropriate one to make welfare judgements is similar to all models which allow for preference changes (see, for example, also O'Donoghue and Rabin, 1999).

4.2 Dictator Game with Initial Entry Decision

Consider the dictator game with entry decision as shown in Figure 6. The game is based on an actual experiment conducted by Dana et al. (2006) in which participants (potential dictators) had to decide whether to enter a 10$ dictator game or not (in that case receiving 9$). In the simplified version considered here, a player first has to decide whether to play a mini dictator game or not. If he refuses he gets 9. If he accepts, he faces the decision of keeping all 10 or to give 3 to another anonymous participant. For the ease of exposition, payoffs for player 2 are dropped. Regarding payoffs at h_1, we assume that the player once actually being in the situation of the dictator is affected by social sharing or fairness norms, reflected in parameters $A, B > 0$ (A capturing potential negative feelings associated with not giving, B capturing potential positive feelings associated with sharing) .

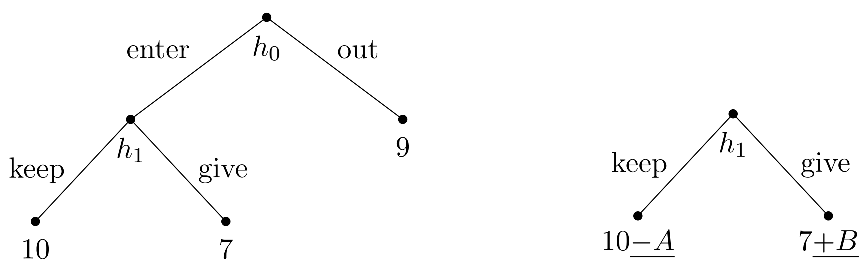

Figure 6: The dictator game with initial entry decision. The left-hand side shows the payoffs as perceived at h_0, while the ones shown on the right-hand side are as perceived upon reaching h_1. Contextual payoffs are underlined.

At h_0 the player knows all the monetary details of the mini dictator game. Given his (emotional) distance to the actual giving decision at history h_1, however, we assume that at h_0 he evaluates the outcomes only in monetary terms.[16] Moreover, with perfect foresight, the player at h_0 knows how his preferences will be different at h_1. Thus, if the contextual influences in h_1 are large enough, i.e. if $A + B > 3$, the player will choose *give* when actually reaching h_1. Foreseeing this, he will choose not to enter the mini dictator game at h_0. In this case, the subgame perfect strategy is (*out give*). Note that, even as presented here, it need not be the case that the participant has no concerns for others per se. What counts is only that he wants to avoid the stimuli of being inside the dictator game which would induce him to give some money (e.g. Andreoni et al., 2017).

Next, consider the case of limited foresight. Again, we assume $\delta = 0$ implying that the player at h_0 believes that his psychological payoffs at h_1, once reached, will not be different. Accordingly, (wrongly) believing he will play *keep* at h_1, the player will enter at h_0 as (*enter keep*) is the unique subgame perfect strategy of the game as perceived at h_0. Once h_1 is reached, though, perceptions change and actual payoffs are as specified on the right-hand

[16] Obviously, the argument can easily be modified in a way that also allows for ex ante positive evaluations of sharing.

side of Figure 6. Accordingly, for A, B large enough, *give* is optimal and (*enter give*) is the unique procedurally justifiable strategy for $\delta = 0$, leaving the player with his least preferred outcome with respect to his preferences at h_0.

Thus, the multistage perspective of the present framework in connection with foresight about future situational effects offers a way to further disentangle possible motivational patterns. For example, according to our model, players who are generally more self-focussed but affected by social norms would be willing to pay to avoid the dictator game, while those who are generally social minded would not (as well as those who do not foresee the effects of social incentives within the dictator situation).

5 Discussion

In the present section, we discuss the relation of our approach to the behavioural literature in more general. In particular, we will try and clarify the analytical possibilities the approach offers by considering foresight about contextual stimuli which may arise at a future point of some interaction. Different from the previous section, the discussion will be more abstract, though, emphasising general patterns rather than addressing specific examples. The section concludes with a rough summary of various common effects that are amenable to an analysis which distinguishes between monetary and socio-psychological payoffs and for which foresight about non-monetary incentives is (arguably) likely to matter.

As we will see, the general pattern of the argument is always roughly the same. As many behavioural models rely on modifications of the standard "rational" utility function, step one is to specify the parts which can be considered as situation specific and which can be treated as (essentially) context independent. What remains is to consider the (general) behavioural consequences that may arise if we consider (limited) foresight about these aspects of utility and do not insist on equilibrium but are content with procedural jusifiability.

Present Biased Preferences

To begin with, consider present biased preferences as discussed by O'Donoghue and Rabin (1999). This model is very close to the present discussion as it already entails multiple periods as well as the possibility of players being more or less aware of (their) future biases. Accordingly, the relation of the models as well as relevance of foresight for the analysis of the corresponding effects are comparably clear.

The observation O'Donoghue and Rabin base their model on is that instantaneous outcomes are usually more relevant to the decision maker than future ones. In their model, they capture this by proposing a utility function which, in addition to a standard monetary aspects and a typical discount factor δ, includes a second discount factor β which is affective between the present and what comes later; the resulting preferences are also known as $(\beta, \delta)-$preferences. For the ease of exposition, we assume δ to be equal to 1 in the sequel (this also avoids confusion with the foresight parameter).[17] The utility function covering time $t \leq T$, as proposed by O'Donoghue and Rabin, then, is given by

$$U'_t(u'_t, u'_{t+1}, \ldots, u'_T) = u'_t + \beta \sum_{\tau=t+1}^{T} u'_\tau,$$

where u'_t represents the utility at time t and $0 < \beta < 1$ incorporates the present bias; the dash superscript is used to avoid confusion with the utility function from our model.

As the situation specific effect considered by O'Donoghue and Rabin (1999) is the additional discounting of future payoffs, a translation of their formulation of utility into the present framework would be the following one:

$$\begin{aligned} u(z,h) &= m(z) + \pi(z,h) \\ &= \underbrace{\sum_{t=0}^{\infty} u'_t(z)}_{\text{'monetary'}} + \underbrace{(\beta - 1) \sum_{t=\tau(h)+1}^{T} u'_t(z)}_{\text{'psychological'}} \end{aligned}$$

where $\tau(h)$ is the history/period in which the decision is actually made.

In the original paper, O'Donoghue and Rabin (1999) discuss two types of biased agents: sophisticates and naifs. The sophisticates know about their present bias and correctly anticipate the preference changes in future periods. Naifs in turn experience the present bias in every stage but always think that the bias is absent in all later periods, i.e. they assume that they later have exactly the same preferences over outcomes as they have now. In the present framework, the two cases would correspond to players having perfect foresight ($\delta = 1$; sophisticates) or no foresight at all ($\delta = 0$; naifs). Moreover, with these specifications, it is straightforward to recover the effects discussed by O'Donoghue and Rabin, namely that sophisticates may act too early doing of unpleasant tasks and too late consuming of temptation goods; and naifs (who's beliefs turn out to be wrong) ending up doing desirable things too early and unpleasant tasks too late. Note that in the original paper there is also a third type of agent, who has no present bias and is called time-consistent. Time-consistents view inter-temporal decisions in a purely monetary way. In our model they correspond to a player, who has no non-zero psychological payoffs.

While the technical similarity of the approaches might seem obvious, it is discussed here to illustrate that for a special case (a situation specific higher valuation of immediate outcomes) the general relevance of phenomena considered in the present paper has already been recognised previously - including the necessity to consider "non-equilibrium" outcomes. Adding to the specific argument provided by O'Donoghue and Rabin, the present more abstract framework essentially allows to study similar patterns also for other context specific effects.

[17] This assumption does not drive results in any significant way and is only made for the sake of expositional ease.

Inequity Aversion

As a further example, consider the model of inequity aversion proposed by Fehr and Schmidt (1999), understood here as one example where preferences for fair behaviour are included in the modelling of utility (the argument is analogous for similar models such as the one proposed by Levine, 1998). In the Fehr-Schmidt model, individuals derive not only utility from economic rewards but also disutility from being behind or ahead of others. While we have no doubt that (many) people do care about fairness, it stands to reason that this is not context independent. For example, Walkers and Wooders (2001) find that professional tennis players are even able to produce what looks like truly random sequences with their service; their argument does not rely on a desire of the players to adjust mutual benefits as much as possible, though. Admittedly, competitive sports is a rather extrem example. Nevertheless, it seems likely that preferences for the well being of others also in general will depend on the context (how close are the others? what is the social surrounding?).

Once we translate the standard Fehr-Schmidt-preferences into the present more abstract formulation of utility, such context effects as well as implications of foresight about these become amenable to a more formal analysis (see also Section 4.2). More specifically, in the n-player form with player i's monetary payoff being denoted by x_i, Fehr-Schmidt utility of player i can be expressed as the sum of his monetary and psychological payoff as follows:

$$U_i'(x;\alpha_i,\beta_i,n) = \underbrace{x_i}_{\text{`monetary'}} \underbrace{- \frac{\alpha_i}{n-1} \cdot \sum_j \max(x_j - x_i, 0) - \frac{\beta_i}{n-1} \cdot \sum_j \max(x_i - x_j, 0)}_{\text{`psychological'}},$$

where concerns for disadvantages are captured by α_i, and for advantages by $\beta_i < \alpha_i$. Note that the empirical evidence of α_i and β_i is indeed mixed (see Dhami, 2016, p. 403, for a broader discussion), which can be viewed as a further indication of the dependence of the effect on the context.

Loss Aversion, Endowment Effect

A further prominent case in point is loss aversion (Kahneman and Tversky, 1979). Here, the effect is that losses appear to have a bigger impact on utility than gains (roughly twice as large). However, what counts as a loss or as a gain is relative to some kind of reference point which in turn is subject to the circumstances of the decision. If we assume that gains do not influence the psychological payoffs while losses come with an additional negative payoff which is of the same size as the monetary loss, we are able to express loss aversion in terms of the proposed utility function:

$$u(z,h) = \underbrace{m(z)}_{\text{`monetary'}} + \underbrace{\min(m(z) - r(h), 0)}_{\text{`psychological'}},$$

with $r(h)$ being the reference point at decision history h. In connection with the multistage set-up considered here and foresight about psychological effects, this again renders possible an analysis of more complex situations where players may (or may not) integrate later changes in reference points in earlier decision making.

In a similar vein, it is possible to express the endowment effect (Thaler, 1980) in terms of the present model. Commonly, the endowment effect – that people tend to value the same good higher if it is in their possession than if it's not – is explained in terms of loss aversion (Kahneman and Tversky, 1979), arguing that selling is perceived as a loss. In the present setting, we can be more agnostic about possible reasons and simply equate the willingness to pay (WTP) with the monetary payoff which can be derived from the consumption of the good. In addition to that, we assume that ownership matters. In particular, if an individual owns the good there is a positive psychological payoff assigned to it (ownership). Accordingly, the seller's willingness to accept (WTA) is given by

$$\text{WTA} = \underbrace{\text{WTP}}_{\text{`monetary'}} + \underbrace{\text{ownership}}_{\text{`psychological'}} .$$

Assuming limited foresight of psychological payoffs, it is immediately obvious that people will underestimate the influence of the endowment effect on both their own and also other people's preferences. Of course, people may learn to incorporate their endowment effect into decision making by observing repeated buying / selling decisions. Yet, this learning experience does not seem to translate into other product categories (cf. Van Boven et al., 2003) so that limited foresight appears to be a plausible assumption.

(Pseudo-)Certainty Effect

The certainty effect goes back to Tversky and Kahneman (1981) and describes the phenomenon that people have a preference for sure outcomes that goes beyond their typical risk preferences.[18] In prospect theory (Kahneman and Tversky, 1979), this effect is attributed to the properties of the probability weighting function. Arguably, assuming a probability weighing function is helpful in terms of the technical understanding and analysis of specific effects. Nonetheless, from an applied point of view, its use is not without problems – not least because the exact form of this function has been subject to some debate (e.g. Ingersoll, 2008; Cavagnaro et al., 2013). At the same time, in most cases, we believe what counts is essentially the fact that small probabilities are special. In the present setting, this can be captured quite easily if we allow outcomes to be lotteries. The preference for sure outcomes then relates to psychological payoffs that are bigger for sure than compared to uncertain outcomes; this formulation encompasses also the pseudo-certainty effect (Tversky and Kah-

[18] According to Tversky and Kahnemann (1981), most people choose a sure gain of $30 compared to getting $45 with a chance of 80% and zero otherwise. Yet, a preference reversal occurs when the probabilities of winning are reduced by a factor of four. A minority chooses the $30 gain which is paid with a probability of 25%, while the majority chooses the $45 gain which is now linked to a realisation probability of 20%.

neman, 1981), which makes outcomes more attractive that are uncertain but perceived as certain. If we denote the effect of a terminal outcome z, being perceived as certain at decision node h, by $c(z,h)$, which is assumed to be increasing in the perceived certainty, utility can be written as follwos:

$$u(z,h) = \underbrace{m(z)}_{\text{`monetary'}} + \underbrace{c(z,h)}_{\text{`psychological'}} .$$

Table 1 summaries the above suggestions about how to express common behavioural effects in a way that makes them amenable to the present framework.

Effect	Corresponding Utility Function	References
Quasi-hyperbolic Discounting	$\underbrace{\sum_{t=0}^{\infty} m_t(z)}_{\text{`monetary'}} + \underbrace{(\beta-1)\sum_{t=\tau(h)+1}^{\infty} m_t(z)}_{\text{`psychological'}}$	O'Donoghue and Rabin (1999)
Interdependent Preferences	$\underbrace{u_i'}_{\text{`monetary'}} + \underbrace{\sum_{j\in N\setminus i} \frac{a_i+\lambda a_j}{1+\lambda} u_j'}_{\text{`psychological'}}, \quad 0 \leq \lambda \leq 1$	Levine (1998)
Inequity Aversion	$\underbrace{x_i}_{\text{`monetary'}} \underbrace{-\frac{\alpha_i}{n-1}\cdot\sum_j \max(x_j-x_i,0) - \frac{\beta_i}{n-1}\cdot\sum_j \max(x_i-x_j,0)}_{\text{`psychological'}}$	Fehr and Schmidt (1999)
Loss Aversion	$\underbrace{m(z)}_{\text{`monetary'}} + \underbrace{\min(m(z)-r(h),0)}_{\text{`psychological'}}$	Tversky and Kahneman (1981)
Endowment Effect	$\text{WTA} = \underbrace{\text{WTP}}_{\text{`monetary'}} + \underbrace{\text{ownership}}_{\text{`psychological'}}$	Thaler (1980)
(Pseudo-)-Certainty Effect	$\underbrace{m(z)}_{\text{`monetary'}} + \underbrace{c(z,h)}_{\text{`psychological'}}$	Tversky and Kahneman (1981)

Table 1: Summary of different behavioural effects and corresponding proposed utility functions expressed so as to fit the current framework. Rows 1-3 follow the original formulas. Rows 4-6 are descriptions presented here. The Table is not comprehensive. Effects mentioned are exemplary.

Projection Bias, Hot-Cold Bias, and Other Effects
Of course, the behavioural effects discussed above are only exemplary. Table 2 below provides a list (again incomplete) of other psychological/behavioural effects which have been considered in the literature and which can be translated into the present framework along similar lines.

A further effect that deserves a special mention at this point as it also emphasises the difficulties in assessing future outcomes is the so called projection bias (e. g. Loewenstein et al. 2003). In brief, the point discussed in that literature is that it is hard to predict ones own future tastes independently from present tastes which are influenced by all kinds of surrounding factors or stimuli; in fact, it can be seen as folk wisdom that you should not go to the supermarket when you are hungry if you want to avoid buying things you later do not like anymore (Gilbert et al., 2002).

Using state-dependent utility Loewenstein et al. (2003) modell the projection bias - the incorrect prediction of future utility - adding a second variable to the utility function, which becomes $u(c, s')$, where c is a period consumption and s' a state, which parameterises the agent's tastes. The prediction about future tastes is denoted by $\tilde{u}(c, s'|s)$, meaning that the agent estimates $u(c, s')$ while he is in state s. The effect of the projection bias, then, can be written as

$$\tilde{u}(c, s'|s) = \beta u(c, s') + (1 - \beta)u(c, s),$$

where β represents the strength of the projection bias ($\beta = 1$ corresponds to no projection bias, $\beta = 0$ to no knowledge of the projection bias / taste changes).

A prominent special case of the projection bias is the hot-cold empathy gap discussed by Loewenstein (2000). It essentially describes the projection bias between a hot state, in which mainly non-monetary aspects count, and a cold state, in which monetary aspects are more important.[19] The mini dictator game from section 4.2 is an example for the reverse version of this effect (the entry decision being made in the emotionally "cold" state).

Obviously, the perceived utility in the present setting is of the same form as the model by Loewenstein et al. (2003), the equivalent to the projection bias parameter β being $\delta^{d(h,h')}$, which depends on the distance between the histories h and h'. The difference in the approaches lies in the fact that we single out a context-independent aspect of utility (related to economic incentives, which we believe to be comparably constant) in order to better differentiate between effects and provide a more formal game theoretic framework which allows to abstractly study the various different corresponding effects – notably without necessitating a commitment to a specific utility function.

[19] Loewenstein explains the empathy gaps by visceral factors, which collect emotions, drive and feeling states. They are able to influence behaviour but are not stable over time. In our model such visceral factors are captured in psychological payoffs and the empathy gap is the result of limited foresight of such psychological payoffs.

Origin	Categories	Description	Literature
The social context / Identity	social norms	Subjects have a generally stable willingness to sacrifice money to take behaviours that are socially appropriate.	Fehr and Fischbacher, 2004; Krupka and Weber, 2013
	social punishment	Cooperators punish free-riders even if it is costly for themselves.	Fehr and Gächter, 2000
	antisocial punishment	Cooperators get punished for their prosocial behaviour.	Herrmann et al., 2008
	audience effect	Cues of observability affect prosocial behaviour.	Haley and Fessler, 2005
	experimenter demand effect	Behaviour by experimental subjects changes due to cues about what constitutes appropriate behaviour.	Zizzo, 2010
	social comparisons	The information about what similar individuals do affects own behaviour.	Gächter et al., 2012
	herd behavior	People mimic behaviour of others even if this means that they have to be ignorant of their (useful) private information.	Scharfstein and Stein, 1990
	reactance	People have a preference for (perceived) freedom.	Kirchler, 1999
	pay what you want	Individuals feel bad when they pay less than the appropriate price, causing them to pass on the opportunity to purchase the product altogether.	Gneezy et al., 2012
	trust	Trustors anticipate that trustees will pay back although this is in terms of money not beneficial for the trustee.	Berg et al., 1995; Kosfeld et al., 2005
	prestige	Giving people the option to report their contributions results in more giving.	Andreoni and Petrie, 2004
	self-esteem	There is a trade-off between money and self-esteem when decision makers have to choose between the two.	Zhang, 2009
	cognitive dissonance	If people face information that contradicts their worldview they get engaged in behaviour to resolve their internal conflict.	Brehm, 1956; Festinger and Carlsmith, 1959
	promises	People have a preference for promise keeping.	Vanberg, 2008
	honesty / lying	People tell the truth even if lying could not be detected and would be beneficial in a monetary sense.	Fischbacher and Föllmi-Heusi, 2013

Origin	Categories	Description	Literature
The context "inside" the player / Visceral states and emotions	temptation	People have a preference for immediate psychological payoffs.	Brown et al., 1996
	anticipatory emotions	*Emotions, that are experienced at the time of decision making, e.g. moods, dread, fear, anxiety.*	
	• moods	Different moods of the same valance may have distinct influences on decision making.	Raghunathan and Pham, 1999; Lerner and Weber, 2013
	anticipated emotions	*Emotions, that are not experienced at the time of decision making, e.g. regret, disappointment, guilt.*	
	• regret	People anticipate regret they might feel later when they realise that they could have made a better decision.	Loomes and Sugden, 1987; Zeelenberg, 1999
	• disappointment	People anticipate disappointment when actual outcomes fall short of expectations.	Bell, 1985; Zeelenberg et al., 1998
	• guilt	People experience guilt if they let down the payoff expectations of others.	Charness and Dufwenberg, 2006; Ellingsen et al., 2010
The context surrounding the decision	anchoring	The information presented first has a big influence on decision making; even when it is not relevant in a rational way.	Tversky and Kahneman, 1974
	framing	The description influences decision making even if the included information is not altered.	Tversky and Kahneman, 1981
	reference points	Reference points are part of the context and determine which outcomes are subject to loss aversion.	Tversky and Kahneman, 1991
	focal points	Players' payoffs in pure coordination games are often better than the resulting outcome of complete randomisation between possible equilibrium strategies.	Mehta et al., 1994
	prominent numbers	Individuals choose prominent numbers disproportionately.	Converse and Dennis, 2018

Table 2: Possible contextual influences on decision making in economic experiments. The selection of literature is only exemplary. More on the respective topics can be found in the cited literature.

6 Concluding Remarks

In this paper, we have proposed a simple framework to model procedural decision making in cases where utility depends on both (foreseeable) monetary and psychological payoffs - the latter being connected to the context of the decision and more difficult to foresee. In doing so, we have used the stylised model of utility proposed by Bergh and Wichardt (2018) which simply treats utility as an additive combination of monetary and context-dependent psychological payoffs. Building on the work by Grossi and Turrini (2012), we have then defined a procedural decision process where every player behaves as a payoff maximiser but may be limited in his ability to foresee all aspects of future payoffs.

In assuming players to be (potentially) biased in their perception of future utility, our argument can be seen as extending earlier ideas by Loewenstein et al. (2003) known as projection bias. Considering also the case of multiple players and dropping standard equilibrium requirements, we have argued that a plausible way to assess behaviour in strategic interactions would be to look for procedurally justifiable strategy profiles, i.e. profiles of behaviour which, at each decision node, are optimal conditional on the perception of the player acting at that node. Moreover, we have demonstrated how assuming psychological payoffs to be subject to limited foresight in the proposed way allows us to cover various well known behavioural inconsistencies in decision making in one abstract formal framework.

As pointed out by Sobel (2005) it is important to find models that allow for preferences to change with the context. With the present argument, we follow his suggestion and include contextual factors in a standard decision making process of a payoff maximising agent. While certainly only a step in that direction, we hope that the arguments provided within the present discussion can help to "identify general properties of extended preferences" as suggested by Sobel (2005, p. 432) and thereby shed some further light onto some still darker parts of this puzzle.

References

Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics, 115*(3), 715-753. https://doi.org/10.1162/003355300554881

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal, 100*(401), 464-477. https://doi.org/10.2307/2234133

Andreoni, J., & Petrie, R. (2004). Public goods experiments without confidentiality: A glimpse into fund-raising. *Journal of Public Economics, 88*(7-8), 1605-1623. https://doi.org/10.1016/S0047-2727(03)00040-9

Andreoni, J., Rao, J. M., & Trachtman, H. (2017). Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. *Journal of Political Economy, 125*(3), 625-653. https://doi.org/10.1086/691703

Apperly, I. (2012). What is "theory of mind"? Concepts, cognitive processes and individual differences. *Quarterly Journal of Experimental Psychology* 65 (5), 825–839. https://doi.org/10.1080/17470218.2012.676055

Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations Research, 33*(1), 1-27. https://doi.org/10.1287/opre.33.1.1

Bentham, J., & Schofield, P. (2008). *A comment on the commentaries and a fragment on government* (Vol. 2). Oxford University Press on Demand.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior, 10*(1), 122-142. https://doi.org/10.1006/game.1995.1027

Bergh, A., & Wichardt, P. C. (2018). Accounting for context: Separating monetary and (uncertain) social incentives. *Journal of Behavioral and Experimental Economics, 72*, 61-66. https://doi.org/10.1016/j.socec.2017.11.002

Bergh, A., & Wichardt, P. C. (2018b). Mine, Ours or Yours? Unintended Framing Effects in Dictator Games. SSRN Working Paper. http://dx.doi.org/10.2139/ssrn.3208589

Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology, 52*(3), 384. https://doi.org/10.1037/h0041006

Brown, K. C., Harlow, W. V., & Starks, L. T. (1996). Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *The Journal of Finance, 51*(1), 85-110. https://doi.org/10.1111/j.1540-6261.1996.tb05203.x

Camerer, C. (2003). *Behavioral Game Theory.* Princeton University Press.

Cavagnaro, D, Pitt, M., Gonzalez, R., & Myung J. (2013). Discriminating among probability weighting functions using adaptive design optimization. *Journal of Risk and Uncertain 47*, 255–289 https://doi.org/10.1007/s11166-013-9179-3

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics, 117*(3), 817-869. https://doi.org/10.1162/003355302760193904

Charness, G., & Dufwenberg, M. (2006). Promises and partnership. *Econometrica, 74*(6), 1579-1601. https://doi.org/10.1111/j.1468-0262.2006.00719.x

Cookson, R. (2000). Framing Effects in Public Goods Experiments. *Experimental Economics* 3, 55–79. https://doi.org/10.1023/A:1009994008166

Converse, B. A., & Dennis, P. J. (2018). The role of "Prominent Numbers" in open numerical judgment: Strained decision makers choose from a limited set of accessible numbers. *Organizational Behavior and Human Decision Processes, 147*, 94-107. https://doi.org/10.1016/j.obhdp.2018.05.007

Dana, J., Cain, D. M., & Dawes, R. M. (2006). What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes, 100*(2), 193-201. https://doi.org/10.1016/j.obhdp.2005.10.001

Dhami, S. (2016). *The foundations of behavioral economic analysis.* Oxford University Press.

Ellingsen, T., Johannesson, M., Tjøtta, S., & Torsvik, G. (2010). Testing guilt aversion. *Games and Economic Behavior, 68*(1), 95-107. https://doi.org/10.1016/j.geb.2009.04.021

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics, 114*(3), 817-868. https://doi.org/10.1162/003355399556151

Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review, 90*(4), 980-994. https://doi.org/10.1257/aer.90.4.980

Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior, 25*(2), 63-87. https://doi.org/10.1016/S1090-5138(04)00005-4

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *The Journal of Abnormal and Social Psychology, 58*(2), 203. https://doi.org/10.1037/h0041593

Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise – an experimental study on cheating. *Journal of the European Economic Association, 11*(3), 525-547. https://doi.org/10.1111/jeea.12014

Gächter, S., Nosenzo, D., & Sefton, M. (2012). The impact of social comparisons on reciprocity. *The Scandinavian Journal of Economics, 114*(4), 1346-1367. https://doi.org/10.1111/j.1467-9442.2012.01730.x

Gilbert, D. T., Gill, M. J., & Wilson, T. D. (2002). The future is now: Temporal correction in affective forecasting. *Organizational Behavior and Human Decision Processes, 88*(1), 430-444. https://doi.org/10.1006/obhd.2001.2982

Gneezy, A., Gneezy, U., Riener, G., & Nelson, L. D. (2012). Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences, 109*(19), 7236-7240. https://doi.org/10.1073/pnas.1120893109

Grossi, D., & Turrini, P. (2012). Short sight in extensive games. In *AAMAS*, 805-812. https://doi.org/10.5555/2343776.2343812

Haley, K. J., & Fessler, D. M. (2005). Nobody's watching?: Subtle cues affect generosity in an anonymous economic game. *Evolution and Human behavior, 26*(3), 245-256. https://doi.org/10.1016/j.evolhumbehav.2005.01.002

Halpern, J. Y., & Rêgo, L. C. (2014). Extensive games with possibly unaware players. *Mathematical Social Sciences, 70*, 42-58. https://doi.org/10.1016/j.mathsocsci.2012.11.002

Heifetz, A., Meier, M., & Schipper, B. C. (2008). A canonical model for interactive unawareness. *Games and Economic Behavior, 62*(1), 304-324. https://doi.org/10.1016/j.geb.2007.07.003

Herrmann, B., Thöni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science, 319*(5868), 1362-1367. https://doi.org/10.1126/science.1153808

Ingersoll, J. (2008). Non-Monotonicity of the Tversky- Kahneman Probability-Weighting Function: A Cautionary Note. *European Financial Management 14*(3), 385–390 https://doi.org/10.1111/j.1468-036X.2007.00439.x

Jehiel, P. (2001). Limited foresight may force cooperation. *The Review of Economic Studies, 68*(2), 369-391. https://doi.org/10.1111/1467-937X.00173

Jehiel, P., & Lilico, A. (2010). Smoking today and stopping tomorrow: a limited foresight perspective. *CESifo Economic Studies, 56*(2), 141-164. https://doi.org/10.1093/cesifo/ifp028

Kahneman, D. (1994). New challenges to the rationality assumption. *Journal of Institutional and Theoretical Economics/ Zeitschrift für die gesamte Staatswissenschaft, 150*(1), 18-36. https://www.jstor.org/stable/40753012

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-292. https://doi.org/10.2307/1914185

Katznelson, H. (2014). Reflective functioning: A review. *Clinical Psychology Review 34* (2), 107-117. https://doi.org/10.1016/j.cpr.2013.12.003

Kirchler, E. (1999). Reactance to taxation: Employers' attitudes towards taxes. *The Journal of Socio-Economics, 28*(2), 131-138. https://doi.org/10.1016/S1053-5357(99)00003-7

Kosfeld, M., Heinrichs, M., Zak, P. J., Fischbacher, U., & Fehr, E. (2005). Oxytocin increases trust in humans. *Nature, 435*(7042), 673-676. https://doi.org/10.1038/nature03701

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic theory, 27*(2), 245-252. https://doi.org/10.1016/0022-0531(82)90029-1

Krupka, E. L., & Weber, R. A. (2013). Identifying social norms using coordination games: Why does dictator game sharing vary?. *Journal of the European Economic Association, 11*(3), 495-524. https://doi.org/10.1111/jeea.12006

Kuhn, H. W. (1953). Extensive games and the problem of information. *Contributions to the Theory of Games, 2*(28), 193-216. https://doi.org/10.1515/9781400881970-012

Lerner, J. S., Li, Y., & Weber, E. U. (2013). The financial costs of sadness. *Psychological Science, 24*(1), 72-79. https://doi.org/10.1177/0956797612450302

Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics, 1*(3), 593-622. https://doi.org/10.1006/redy.1998.0023

Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *American Economic Review, 90*(2), 426-432. https://doi.org/10.1257/aer.90.2.426

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin, 127*(2), 267. https://doi.org/10.1037/0033-2909.127.2.267

Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *The Quarterly Journal of Economics, 118*(4), 1209-1248. https://doi.org/10.1162/003355303322552784

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal, 92*(368), 805-824. https://doi.org/10.2307/2232669

Loomes, G., & Sugden, R. (1987). Testing for regret and disappointment in choice under uncertainty. *The Economic Journal, 97*, 118-129. https://doi.org/10.2307/3038234

Mehta, J., Starmer, C., & Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review, 84*(3), 658-673. https://www.jstor.org/stable/2118074

Mischel, W., Shoda, Y., & Rodriguez, M. I. (1989). Delay of gratification in children. *Science, 244*(4907), 933-938. https://doi.org/10.1126/science.2658056

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review, 89*(1), 103-124. https://doi.org/10.1257/aer.89.1.103

Osborne, M. J., & Rubinstein, A. (1994). *A course in game theory.* MIT Press.

Perea, A. (2014). Belief in the opponents' future rationality. *Games and Economic Behavior, 83*, 231-254. https://doi.org/10.1016/j.geb.2013.11.008

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review, 83*(5), 1281-1302. https://www.jstor.org/stable/2117561

Raghunathan, R., & Pham, M. T. (1999). All negative moods are not equal: Motivational influences of anxiety and sadness on decision making. *Organizational Behavior and Human Decision Processes, 79*(1), 56-77. https://doi.org/10.1006/obhd.1999.2838

Scharfstein, D. S., & Stein, J. C. (1990). Herd behavior and investment. *American Economic Review, 80*(3), 465-479. https://www.jstor.org/stable/2006678

Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil i: Bestimmung des dynamischen Preisgleichgewichts. *Journal of Institutional and Theoretical Economics/ Zeitschrift für die gesamte Staatswissenschaft, 121*(2), 301-324. https://www.jstor.org/stable/40748884

Shackleton, R. (1993). The greatest happiness of the greatest number: The history of Bentham's phrase. In *Jeremy Bentham: Critical Assessments, 2*, (pp. 353-366).

Shoda, Y., Mischel, W., & Peake, P. K. (1990). Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions. *Developmental Psychology, 26*(6), 978. https://doi.org/10.1037/0012-1649.26.6.978

Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of Economic Literature, 43*(2), 392-436. https://doi.org/10.1257/0022051054661530

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies, 23*(3), 165-180. https://doi.org/10.2307/2295722

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization, 1*(1), 39-60. https://doi.org/10.1016/0167-2681(80)90051-7

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124-1131. https://doi.org/10.1126/science.185.4157.1124

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*(4481), 453-458. https://doi.org/10.1126/science.7455683

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics, 106*(4), 1039-1061. https://doi.org/10.2307/2937956

Vanberg, C. (2008). Why do people keep their promises? An experimental test of two explanations. *Econometrica, 76*(6), 1467-1480. https://doi.org/10.3982/ECTA7673

Van Boven, L., Loewenstein, G., & Dunning, D. (2003). Mispredicting the endowment effect:: Underestimation of owners' selling prices by buyer's agents. *Journal of Economic Behavior & Organization, 51*(3), 351-365. https://doi.org/10.1016/S0167-2681(02)00150-6

Van Boven, L., & Loewenstein, G. (2005). Empathy gaps in emotional perspective taking. In *Other Minds: How Humans Bridge the Divide between Self and Others*, 284-297. Guilford Press. https://www.andrew.cmu.edu/user/gl20/GeorgeLoewenstein /Papers_files/pdf/EmpathyGapsEmotPersp.pdf

Walker, M., & Wooders, J. (2001), Minimax Play at Wimbledon. *American Economic Review* 91 (5), pp. 1521-1538. https://doi: 10.1257/aer.91.5.1521

Watts, T. W., Duncan, G. J., & Quan, H. (2018). Revisiting the marshmallow test: A conceptual replication investigating links between early delay of gratification and later outcomes. *Psychological Science, 29*(7), 1159-1177. https://doi.org/10.1177/0956797618761661

Wichardt, P. C. (2008). Identity and why we cooperate with those we do. *Journal of Economic Psychology, 29*(2), 127-139. https://doi.org/10.1016/j.joep.2007.04.001

Wichardt, P. C. (2010). Modelling equilibrium play as governed by analogy and limited foresight. *Games and Economic Behavior, 70*(2), 472-487. https://doi.org/10.1016/j.geb.2010.03.002

Zeelenberg, M., van Dijk, W. W., Van der Pligt, J., Manstead, A. S., Van Empelen, P., & Reinderman, D. (1998). Emotional reactions to the outcomes of decisions: The role of counterfactual thought in the experience of regret and disappointment. *Organizational Behavior and Human Decision Processes, 75*(2), 117-141. https://doi.org/10.1006/obhd.1998.2784

Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making, 12*(2), 93-106. https://doi.org/10.1002/(SICI)1099-0771(199906)12:2<93::AID-BDM311>3.0.CO;2-S

Zhang, L. (2009). An exchange theory of money and self-esteem in decision making. *Review of General Psychology, 13*(1), 66-76. https://doi.org/10.1037/a0014225

Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics, 13*(1), 75-98. https://doi.org/10.1007/s10683-009-9230-z