

Frustration and Anger in Games*

Pierpaolo Battigalli[†] Martin Dufwenberg[‡] Alec Smith[§]

February 23, 2015

Abstract

Frustration, anger, and aggression have important consequences for economic and social behavior, concerning for example monopoly pricing, contracting, bargaining, traffic safety, violence, and politics. Drawing on insights from psychology, we develop a formal approach to exploring how frustration and anger, via blame and aggression, shape interaction and outcomes in economic settings.

KEYWORDS: frustration, anger, blame, belief-dependent preferences, psychological games

JEL codes: C72, D03

*This paper modifies and extends Smith (2009). Pierpaolo Battigalli gratefully acknowledges financial support from ERC advanced grant 324219. We thank Paolo Leonetti and Marco Stenborg Petterson for excellent research assistance, Pierfrancesco Guarino and multiple seminar audiences for helpful comments.

[†]Bocconi University and IGIER. email: pierpaolo.battigalli@unibocconi.it.

[‡]Bocconi University, IGIER, University of Arizona, University of Gothenburg, CESifo. email: martin.dufwenberg@unibocconi.it.

[§]University of Arizona and Compass Lexecon. email: alecs@email.arizona.edu.

1 Introduction

Anger may play a key role for shaping outcomes in economically important ways. Consider three cases:

Case 1: In 2006 US gas prices went up and up. Many folks were upset. Did this cause road rage, or people trading a truck for a Hyundai? Did gas stations or truck dealers go easier on price hikes, or offer rebates, anticipating potential adverse effects on sales that might otherwise materialize?

Case 2: When local football teams that are favored to win instead lose, the police get more reports of husbands assaulting wives (Card & Dahl 2011). Do unexpected losses spur thus vented frustration?

Case 3: Following Sovereign Debt Crises (2009-), some EU countries embarked on austerity programs. Was it because citizens lost benefits that some cities experienced riots?

Traffic safety, pricing, domestic violence, political landscapes: the examples above illustrate some situations where anger may have important consequences. However, to carefully assess how emotions such as anger may shape social and economic interactions, one needs a theory that predicts outcomes based on the decision-making of anger-prone individuals and that also accounts for the strategic consideration of such individuals' behavior by their co-players. Our paper develops such a theory.

Insights from psychology about both the triggers of anger and its consequences for behavior suggest how to incorporate anger into models of strategic interaction. The behavioral consequences of emotions are referred to as "action tendencies," and the action tendency associated with anger is aggression. One may imagine that angry players are willing to forego material gains to punish others, or that a predisposition to behave aggressively when angered may benefit a player by serving as a credible threat, and so on. But while insights of this nature can be gleaned from psychologists' writings, their analysis usually stops with the individual rather than going on to assess overall economic/social implications. We take the basic insights about anger that psychology has produced as input and inspiration for the theory

we develop and apply.¹

Economists have traditionally paid scant attention to emotions, including anger, but interest is on the rise. Several recent studies inspire us. Most are empirical, indicative of hostile action occurring in economic situations, based on either observational² or experimental³ data. A few of these studies present theory,⁴ typically with the purpose of explaining specific data patterns. Our approach is different. We do not start with data, but with notions from psychology which we incorporate in general games. We are led to models that differ substantially from the existing theory, though predictions may be similar in their specific settings.

Psychologists suggest that anger is typically anchored in frustration, which occurs when someone is unexpectedly denied something he or she cares about.⁵ We assume (admittedly restrictively; cf. section 7) that people are frustrated when they get less material rewards than they expected beforehand. Moreover, they then become hostile towards whomever they blame. There are several ways that blame may be assigned (cf. Alicke 2000) and we present three distinct approaches, captured by distinct utility functions. While players motivated by simple anger (SA) become generally hostile when frustrated, those motivated by anger from blaming behavior (ABB) or by anger from blaming intentions (ABI) go after others more discriminately, asking who caused, or who intended to cause, their dismay.

What are the overall implications when people interact? To provide answers, we develop a notion of polymorphic sequential equilibrium (PSE). Players are assumed to correctly anticipate how others behave on average, and the concept furthermore allows for different “types” of the same player to have different plans in equilibrium, which yields meaningful updating of

¹The relevant literature is huge. A good point of entry, and source of insights and inspiration for us, is the recent *International Handbook of Anger* (Potegal, Spielberger & Stemmler 2010), which offers a cross-disciplinary perspective over 32 chapters reflecting “affective neuroscience, business administration, epidemiology, health science, linguistics, political science, psychology, psychophysiology, and sociology” (p. 3, opening chapter). We take the non-occurrence of “economics” in the list as an indication our approach is original and needed!

²See Anderson & Simester (2010) and Rotemberg (2005, 2011) on pricing; Card & Dahl on domestic violence.

³See Carpenter & Matthews (2012), Gurdal, Miller & Rustichini (2014), Gneezy & Imas (2014).

⁴See Rotemberg (2005, 2008, 2011), Akerlof (2013), Passerelli & Tabellini (2013)

⁵Psychologists often refer to this as “goal-blockage;” cf. p.3 of the (*op.cit.*) *Handbook*.

players' views of others' intentions as various subgames are reached. This is crucial for a sensible treatment of how players consider intentionality as they blame others. We apply this solution concept to the aforementioned utility functions, explore properties, and compare predictions.

A player's frustration depends on his beliefs about others' choices. The blame a player attributes to another may depend on his beliefs about others' choices or beliefs. For these reasons, all our models find their intellectual home in the framework of psychological game theory; see Geanakoplos, Pearce & Stacchetti (1989), Battigalli & Dufwenberg (2009).

We develop most of our analysis within a two-period setting described in Section 2. Section 3 defines frustration. Section 4 develops our three key notions of psychological utility. Section 5 introduces the equilibrium concept and derives/highlights various results and insights. Section 6 generalizes the analysis to multistage games. Section 7 concludes.

2 Setup

Players engage in a two-stage interaction. Stage, or period $t \in \{1, 2\}$ is the time interval between dates $t - 1$ and t . The set of active players and their feasible actions depend on the period and on previous choices. Players start with initial beliefs at date zero, and revise their beliefs conditioning on what they learn. We first describe the rules of interaction, or game form, and then we define initial and conditional beliefs.

2.1 Game form

We consider a finite two-stage game form describing the rules of interaction and the consequences of players' actions. The set of players, possibly including passive individuals, is I . Letting \emptyset denote the empty history (the root of the game), there is a finite set $I(\emptyset)$ of first movers; each $i \in I(\emptyset)$ picks an action a_i^1 in the finite feasible set $A_i(\emptyset)$. More generally, $I(h)$ and $A_i(h)$ respectively denote the set of active players and the set of feasible actions of player i at a given history h .⁶ The finite set of active players in the second period,

⁶Formally, it is convenient to adopt the convention that inactive players have only one feasible action, the pseudo-action "wait." Thus $|A_j(h)| = 1$ ($A_j(h)$ is a singleton) whenever $j \notin I(h)$, and $|A_{-i}(h)| = 1$ if $I(h) = \{i\}$. Whenever a set Y is a singleton, we identify the Cartesian product $X \times Y$ with X , as the two sets are isomorphic.

$I(a^1)$, depends on the first-period action profile $a^1 = (a_i^1)_{i \in I(\emptyset)}$, which becomes public information at the beginning of the second period. If $I(a^1) = \emptyset$, the game ends; otherwise, each player $i \in I(a^1)$ chooses an action a_i^2 in the finite feasible set $A_i(a^1)$; the resulting action profile is $a^2 = (a_i^2)_{i \in I(a^1)}$. In each period, active players move simultaneously. We let $A(h)$ denote the set feasible action profiles at h . Similarly, for each player i , $A_{-i}(h)$ denotes the set of feasible action profiles of the co-players at h .⁷

The root \emptyset and the feasible histories $a^1 \in A(\emptyset)$ and $(a^1, a^2) \in A(\emptyset) \times A(a^1)$ are the nodes of the game tree. We let Z denote the set of terminal histories/nodes of the game tree, and H denote the set of non-terminal, or partial histories.⁸ For each $h \in H$, $Z(h)$ denotes the set of terminal successors of h .⁹ If the set of active players $I(h)$ is a singleton for each $h \in H$, that is, if players never move simultaneously, then the game has perfect information.

We assume that the material consequences of players' actions are determined by the profile of payoff functions $(\pi_i : Z \rightarrow \mathbb{R})_{i \in I}$. This completes the description of the game form, if there are no chance moves. If the game contains chance moves, we augment the player set with a dummy player c (with $c \notin I$), who selects a feasible action at random. Thus, we consider an augmented player set $I_c = I \cup \{c\}$, and the sets of first movers and second movers may include c as well: $I(\emptyset), I(a^1) \subseteq I_c$. If the chance player is active at history $h \in H$, the chance move is described by a probability density

⁷According to our convention, $A(h) = \times_{i \in I} A_i(h)$ and $A_{-i}(h) = \times_{j \in I \setminus \{i\}} A_j(h)$, which are – respectively – isomorphic to $\times_{i \in I(h)} A_i(h)$ and $\times_{j \in I(h) \setminus \{i\}} A_j(h)$.

⁸Formally, let

$$Z_1 = \{a^1 : a^1 \in A(\emptyset), I(a^1) = \emptyset\}.$$

Then

$$H = \{\emptyset\} \cup A(\emptyset) \setminus Z_1,$$

$$Z = Z_1 \cup \{(a^1, a^2) : a^1 \in A(\emptyset) \setminus Z_1, a^2 \in A(a^1)\}.$$

⁹That is, $Z(\emptyset) = Z$ and $Z(\bar{a}^1) = \{(a^1, a^2) : a^1 = \bar{a}^1, a^2 \in A(\bar{a}^1)\}$ for each $\bar{a}^1 \in A(\emptyset) \setminus Z_1$.

function (pdf) $\sigma_c(\cdot|h) \in \Delta(A_c(h))$. Table 1 summarizes:

Notation	Terminology
$i \in I$	players
$c, I_c = I \cup \{c\}$	chance, set of players including chance
$t \in \{1, 2\}$	stages, or periods
a_i^t	action of i in stage t
$a^t (a_{-i}^t)$	action profile (of others) in stage t
$h \in H$	non-terminal, or partial histories
$I(h) \subseteq I_c$	set of active players at h
$A_i(h), A(h), A_{-i}(h)$	set of actions and action profiles at h
$\sigma_c(a_c^t h)$	probabilities of chance moves
$z \in Z$	terminal histories
$Z(h)$	terminal successors of h
$\pi_i : Z \rightarrow \mathbb{R}$	monetary payoff function of $i \in I$

Table 1. Elements of the two-stage game form.

The following example illustrates our notation.




Figure A. Asymmetric punishment.

Example 1 (*Asymmetric Punishment*) Ann, Bob and Penny the punisher play the following game form, where Ann and Bob move simultaneously in

the first stage. Penny may move in the second stage; by choosing P she can then decrease π_b (while π_a increases). See Figure A, where profiles of actions and of monetary payoffs are listed according to players' alphabetical order. Using our notation, we have:

$$\begin{aligned} H &= \{\emptyset, (D, L)\}, \\ Z &= \{(U, L), (U, R), (D, L), ((D, L), N), ((D, L), P)\}, \\ I(\emptyset) &= \{a, b\}, I((D, L)) = \{p\}, \\ A_a(\emptyset) &= \{U, D\}, A_b(\emptyset) = \{L, R\}, A_p((D, L)) = \{N, P\}. \quad \blacktriangle \end{aligned}$$

2.2 Beliefs

It is conceptually useful to distinguish the following three aspects of a player's beliefs: beliefs about co-players' actions, beliefs about co-players' beliefs, and the player's plan, which we represent as beliefs about own actions. Beliefs are defined conditional on each history. Let us abstractly denote by Δ_{-i} the space of co-players' beliefs. Player i 's beliefs can be compactly described as conditional probability measures over paths and beliefs of others, that is, over $Z \times \Delta_{-i}$. Events, from i 's point of view, are subsets of $Z \times \Delta_{-i}$. Events about behavior have the form $Y \times \Delta_{-i}$, with $Y \subseteq Z$; events about beliefs have the form $Z \times E_{\Delta_{-i}}$, with $E_{\Delta_{-i}} \subseteq \Delta_{-i}$.¹⁰

Personal histories To model how i determines the subjective value of each feasible action at each history where he is active, we add to the commonly observed histories $h \in H$ also personal histories of the form (h, a_i) , with $i \in I(h)$, $a_i \in A_i(h)$. In a game with perfect information, $(h, a_i) \in H \cup Z$. But if there are simultaneous moves at h , then (h, a_i) is not a history in the standard sense. As soon as i irreversibly chooses action a_i , he observes (h, a_i) , and determines the value of a_i using his beliefs conditional on this event. We denote by H_i the set of histories of i – the standard and personal ones – and by $Z(h, a_i)$ the set of terminal successors of personal history (h, a_i) .¹¹ The standard precedence relation \prec for histories in $H \cup Z$ is extended to H_i in the

¹⁰We assume that Δ_{-i} is a compact metrizable space, which is justified by the construction of hierarchical belief spaces given below. Events are Borel measurable subsets of $Z \times \Delta_{-i}$. We do not specify the terminal beliefs of i about the beliefs of others, because they are not relevant for the models in this paper.

¹¹That is,

$$H_i = H \cup \{(h, a_i) : h \in H, i \in I(h), a_i \in A_i(h)\}$$

obvious way: for all $h \in H$, $i \in I(h)$, and $a_i \in A_i(h)$, it holds that $h \prec (h, a_i)$ and $(h, a_i) \prec (h, (a_i, a_{-i}))$ if i is not the only active player at h .

Conditional probability systems Player i 's system of beliefs β_i is an array of conditional beliefs indexed by histories in H_i : $\beta_i = (\beta_i(\cdot|[h_i]))_{h_i \in H_i}$, where $\beta_i(\cdot|[h_i])$ is a probability measure concentrated on event about behavior $[h_i] = Z(h_i) \times \Delta_{-i}$ for all $h_i \in H_i$. We use obvious abbreviations like

$$\beta_i(h'_i|h_i) = \beta_i([h'_i]|[h_i])$$

whenever this causes no confusion. More generally, we suppress parentheses when this does not compromise understanding.

The **first-order** belief system of i gives the probabilities of terminal histories and of action profiles conditional on each history:

$$\alpha_i(z|h_i) = \beta_i(z|h_i), \alpha_i(a|h) = \alpha_i((h, a)|h) \quad (1)$$

for all $z \in Z$, $h_i \in H_i$, $h \in H$ and $a \in A(h)$.

A belief system β_i must satisfy some natural properties. First of all, the rules of conditional probabilities must hold whenever possible: if $h_i \prec h'_i$ then

$$\beta_i(h'_i|h_i) > 0 \Rightarrow \beta_i(E|h'_i) = \frac{\beta_i(E \cap [h'_i]|h_i)}{\beta_i(h'_i|h_i)} \quad (2)$$

for all $h_i, h'_i \in H_i$ and every event $E \subseteq Z \times \Delta_{-i}$. Equations (1)-(2) imply

$$\alpha_i(a^1, a^2|\emptyset) = \alpha_i(a^2|a^1) \alpha_i(a^1|\emptyset).$$

Second, i realizes that his choice cannot influence simultaneous choices and beliefs of co-players, so i 's beliefs satisfy a causal independence property:

$$\beta_i([h, (a_i, a_{-i})] \cap E_{-i}|(h, a_i)) = \beta_i([h, (a'_i, a_{-i})] \cap E_{-i}|(h, a'_i)), \quad (3)$$

for every $h \in H$, $a_i, a'_i \in A_i(h)$, $a_{-i} \in A_{-i}(h)$, and $E_{-i} = Z \times E_{\Delta_{-i}}$.

Properties (1)-(3) imply

$$\alpha_i(a_i, a_{-i}|h) = \alpha_{i,i}(a_i|h) \alpha_{i,-i}(a_{-i}|h),$$

and

$$Z(h, a_i) = \bigcup_{a_{-i} \in A_{-i}(h)} Z(h, (a_i, a_{-i})).$$

where $\alpha_{i,i}(\cdot|h)$ and $\alpha_{i,-i}(\cdot|h)$ are marginals of $\alpha_i(\cdot|h)$ on $A_i(h)$ and $A_{-i}(h)$.

Note that the array of conditional probabilities $\alpha_{i,i} = (\alpha_{i,i}(\cdot|h))_{h \in H} \in \times_{h \in H} \Delta(A_i(h))$ is – technically speaking – a behavioral strategy, and we interpret it as the **plan** of i . The reason is that the result of i 's contingent planning is precisely a system of conditional beliefs about what action he would take at each history. If there is only one co-player, also $\alpha_{i,-i} \in \times_{h \in H} \Delta(A_{-i}(h))$ formally corresponds to a behavioral strategy. With multiple co-players, $\alpha_{i,-i}$ corresponds instead to a “correlated behavioral strategy.” Whatever the case, $\alpha_{i,-i}$ gives the conditional beliefs of i about the behavior of others, and these beliefs may not coincide with the plans of others. We emphasize that the plan of a player is not an actual choice: actions on the path of play are the only actual choices.

A belief system β_i satisfying (2)-(3) is a **conditional probability system**, or CPS. The set of such CPSs is denoted $\Delta^{H_i}(Z \times \Delta_{-i})$, a subset of $[\Delta(Z \times \Delta_{-i})]^{H_i}$. Whenever this causes no confusion, we write initial beliefs omitting the empty history, as in $\beta_i(E) = \beta_i(E|\emptyset)$, or $\alpha_i(a) = \alpha_i(a|\emptyset)$.

Hierarchical beliefs Of course, (2)-(3) can be directly stated for the system of first-order beliefs $\alpha_i = (\alpha_i(\cdot|h))_{h \in H_i}$, that is, the conditional beliefs about paths. The set of first-order CPSs, $\Delta_i^1 = \Delta^{H_i}(Z)$, is a compact metrizable space. The set of second-order CPSs, $\Delta_i^2 = \Delta^{H_i}(Z \times \Delta_{-i}^1)$ where $\Delta_{-i}^1 = \times_{j \in I \setminus \{i\}} \Delta_j^1$, is compact and metrizable as well.¹² Higher-order belief spaces can be defined by recursion, but we do not need them in our analysis. Note that the first-order CPS $\alpha_i \in \Delta_i^1 = \Delta^{H_i}(Z)$ has to be derived from the second-order CPS $\beta_i \in \Delta_i^2 = \Delta^{H_i}(Z \times \Delta_{-i}^1)$, otherwise i 's second-order hierarchy (α_i, β_i) would be incoherent. Indeed, it can be checked that starting from $\beta_i \in \Delta_i^2$ and letting

$$\alpha_i(Y|h) = \beta_i(Y \times \Delta_{-i}^1|h)$$

for all $h \in H_i$ and $Y \subseteq Z$, we obtain an array $\alpha_i = (\alpha_i(\cdot|h))_{h \in H_i}$ of conditional probabilities satisfying (1)-(3), that is, an element of Δ_i^1 . Whenever we write in a formula beliefs of different orders for the same player, we assume that first-order beliefs are derived from second-order beliefs.

¹²This holds for higher-order beliefs in general, as can be shown following a proof by Battigalli & Siniscalchi (1999) with a minor modification to take independence into account.

Conditional expectations Let ψ_i be any real-valued measurable function of variables that player i does not know, *e.g.*, the terminal history or the co-players' first-order beliefs. Then i can compute the expected value of ψ_i conditional on any common or personal history $h_i \in H_i$ by means of his CPS β_i . This expected value is denoted by $\mathbb{E}[\psi_i|h_i; \beta_i]$. If ψ_i depends only on actions, that is, on the path z , then $\mathbb{E}[\psi_i|h_i; \beta_i]$ is determined by the first-order CPS α_i derived from β_i , and we can write $\mathbb{E}[\psi_i|h_i; \alpha_i]$. In particular, the first-order CPS α_i gives the conditional expected material payoffs:

$$\begin{aligned}\mathbb{E}[\pi_i|h; \alpha_i] &= \sum_{z \in Z(h)} \alpha_i(z|h) \pi_i(z), \\ \mathbb{E}[\pi_i|(h, a_i); \alpha_i] &= \sum_{z \in Z(h, a_i)} \alpha_i(z|h, a_i) \pi_i(z)\end{aligned}$$

for all $h \in H$, $a_i \in A_i(h)$. $\mathbb{E}[\pi_i|h; \alpha_i]$ is what i expects to get conditional on h given CPS α_i , which also specifies i 's plan; while $\mathbb{E}[\pi_i|(h, a_i); \alpha_i]$ is the expected payoff of taking action a_i . If a_i is precisely the action that i planned to take at h , $\alpha_{i,i}(a_i|h) = 1$, then $\mathbb{E}[\pi_i|h; \alpha_i] = \mathbb{E}[\pi_i|(h, a_i); \alpha_i]$. For initial beliefs, we omit $h = \emptyset$ from such expressions; in particular, the initially expected payoff is $\mathbb{E}[\pi_i; \alpha_i]$.

3 Frustration

We will present several models of how frustrated players attribute blame and go after others, but keep our account of frustration constant. Here is the key definition: i 's **frustration** in stage 2, given a^1 , is

$$F_i(a^1; \alpha_i) = \left[\mathbb{E}[\pi_i; \alpha_i] - \max_{a_i^2 \in A_i(a^1)} \mathbb{E}[\pi_i|(a^1, a_i^2); \alpha_i] \right]^+,$$

where $[x]^+ = \max\{x, 0\}$. Player i 's frustration in stage 2 is given by the gap, if positive, between his initially expected payoff and the currently best expected payoff he believes he can obtain. Diminished expectation $-\mathbb{E}[\pi_i|a^1; \alpha_i] < \mathbb{E}[\pi_i; \alpha_i]$ – is only a necessary condition for frustration. For i to be frustrated it must also be the case that i cannot close the gap. This captures the psychological intuition that, given i 's beliefs, i 's frustration in stage 2 does not depend on his stage 2 action. Had we alternatively modeled frustration

as equal to actual diminished expectations (that is, $\mathbb{E}[\pi_i; \alpha_i] - \mathbb{E}[\pi_i|a^1; \alpha_i]$), this would have had counterintuitive implications.

$F_i(a^1; \alpha_i)$ expresses stage 2 frustration. One could define frustration at the root, or at end nodes, but neither would matter for our purposes. At the root nothing has happened, so frustration equals zero. Frustration is possible at the end nodes, but can't influence subsequent choices as the game is over. One might allow the anticipated frustration at end nodes to influence earlier decisions; however, to simplify the assumptions we make in the analysis below, we rule out this possibility. Furthermore, players are influenced by the frustrations of co-players only to the extent that they affect behavior.

To illustrate the definition of $F_i(a^1; \alpha_i)$, return to the game form of Example 1:

Example 2 Suppose that, in the game form of Figure A, Penny initially expects to get \$2, i.e., $\alpha_p((U, L)|\emptyset) + \alpha_p((D, R)|\emptyset) = 1$ and $\mathbb{E}[\pi_p; \alpha_p] = 2$. Penny's frustration after $a^1 = (D, L)$ is

$$F_p((D, L); \alpha_i) = [\mathbb{E}[\pi_p; \alpha_p] - \max\{\pi_p((D, L), N), \pi_p((D, L), P)\}]^+ = 2 - 1 = 1.$$

Penny's frustration is independent of her plan, because she is initially certain she will not move. Suppose instead that $\alpha_p((U, L)|\emptyset) = \alpha_p((D, L)|\emptyset) = \frac{1}{2}$. Then

$$F_p((D, L); \alpha_i) = \frac{1}{2} \times 2 + \frac{1}{2} \alpha_p(N|(D, L)) \times 1 - 1 = \frac{1}{2} \alpha_p(N|(D, L)).$$

Penny's frustration is now highest when she initially plans not to punish Bob. However, her frustration after a^1 is independent of her actual choice: Penny's frustration equals $\frac{1}{2} \alpha_p(N|(D, L))$ independently of whether she ultimately chooses N or P . \blacktriangle

4 Anger

A player's preferences over actions at a given node – his action tendencies – depend on expected material payoffs and frustration. A frustrated player tends to hurt others, if this is not too costly (cf. Dollard *et al.* 1939, Averill 1983, Berkowitz 1989). We consider different versions of this frustration-aggression hypothesis related to different levels of cognitive appraisal. In

general, player i moving at history h chooses action a_i in order to maximize the expected value of a belief-dependent “decision utility” of the form

$$u_i(h, a_i; \beta_i) = \mathbb{E}[\pi_i | (h, a_i); \alpha_i] - \theta_i \sum_{j \neq i} B_{ij}(h; \beta_i) \mathbb{E}[\pi_j | (h, a_i); \alpha_j], \quad (4)$$

where $B_{ij}(h; \beta_i) \geq 0$ measures how much of i 's frustration is blamed on co-player j (and hence the tendency to hurt j), α_i is the first-order CPS derived from second-order belief β_i , and θ_i is a sensitivity parameter. We assume that $B_{ij}(h; \beta_i)$ is positive only if frustration is positive:

$$B_{ij}(a^1; \beta_i) \leq F_i(a^1; \alpha_i). \quad (5)$$

Therefore, the decision utility of a first-mover coincides with expected material payoff, because there cannot be any frustration in the first stage:

$$u_i(\emptyset, a_i; \beta_i) = \mathbb{E}[\pi_i | a_i; \alpha_i].$$

When i is the only active player at $h = a^1$, he determines the terminal history with his choice $a_i = a^2$, and decision utility has the form

$$u_i(a^1, a_i; \beta_i) = \pi_i(a^1, a_i) - \theta_i \sum_{j \neq i} B_{ij}(a^1; \beta_i) \pi_j(a^1, a_i).$$

We now proceed to consider three specific functional forms that capture different notions of blame.

4.1 Simple Anger (SA)

Our most rudimentary hypothesis, **simple anger (SA)**, is that i 's tendency to hurt others is proportional to i 's frustration, un-modulated by cognitive appraisal of blame, so $B_{ij}(a^1; \beta_i) = F_i(a^1; \alpha_i)$:

$$u_i^{SA}(a^1, a_i; \alpha_i) = \mathbb{E}[\pi_i | (a^1, a_i); \alpha_i] - \theta_i \sum_{j \neq i} F_i(a^1; \alpha_i) \mathbb{E}[\pi_j | (a^1, a_i); \alpha_j]. \quad (6)$$




Figure B. Ultimatum Minigame.

Example 3 (*Ultimatum Minigame*) Ann and Bob play a simple bargaining game: Ann can make a fair offer, which is automatically accepted, or a greedy offer which Bob can either accept or reject. See Figure B. Bob’s frustration following the greedy offer is

$$F_b(g; \alpha_b) = [(2(1 - \alpha_b(g)) + \alpha_b(g)\alpha_b(y)) - 1]^+.$$

Therefore

$$u_b^{SA}(g, n; \alpha_i) - u_b^{SA}(g, y; \alpha_i) = 3\theta_b [(2(1 - \alpha_b(g)) + \alpha_b(g)\alpha_b(y)) - 1]^+ - 1.$$

For Bob to be frustrated he must not expect the greedy offer with certainty. If he is frustrated, the less he expects the greedy offer, and – interestingly – the less he plans to reject it, the more prone he is to reject once the greedy offer materializes. The more resigned Bob is to getting a low payoff, the less frustrated and prone to aggression he will be when receiving the low-ball offer. ▲

4.2 Anger from blaming behavior (ABB)

Action tendencies may depend on a player’s cognitive appraisal of how to blame others. When a frustrated player i blames co-players for their behavior, he looks only at the actions chosen in stage 1, without considering intentions, that is, without considering others’ plans and beliefs about others. Player i ’s tendency to hurt j is determined by a continuous blame function $B_{ij}(a^1; \alpha_i)$ that depends only on first-order belief α_i such that

$$B_{ij}(a^1; \alpha_i) = \begin{cases} 0, & \text{if } j \notin I(\emptyset), \\ F_i(a^1; \alpha_i), & \text{if } \{j\} = I(\emptyset). \end{cases} \quad (7)$$

Equation (7) says that, if j is not active in the first stage, he cannot be blamed for i 's frustration, and if instead j is the *only* active player, he is fully blamed.¹³ We consider below specific functional forms for $B_{ij}(a^1; \alpha_i)$ that satisfy (5)-(7). With this, the decision utility with **anger from blaming behavior (ABB)** is

$$u_i^{ABB}(a^1, a_i; \alpha_i) = \mathbb{E} [\pi_i | (a^1, a_i); \alpha_i] - \theta_i \sum_{j \neq i} B_{ij}(a^1; \alpha_i) \mathbb{E} [\pi_j | (a^1, a_i); \alpha_i].$$

The following example illustrates the difference between SA and ABB:




Figure C. Hammering one's thumb.

Example 4 (*Inspired by Frijda, 1993*) Andy the handyman (player a in Figure C) uses a hammer. His apprentice, Bob, has no payoff-relevant action. In a bad day (determined by chance) Andy hammers his thumb and can then either take it out on Bob or not. If he does, he further disrupts production. See Figure C. Assuming $\alpha_a(B) = \varepsilon < 1/2$, the extent of Andy's frustration in a bad day is

$$F_a(B; \alpha_a) = 2(1 - \varepsilon) + \varepsilon \alpha_a(N|B) - 1 > 0.$$

With SA and θ_a sufficiently high, Andy takes it out on Bob. But, since Bob is passive, with ABB Andy chooses N regardless of θ_a . **▲**

SA and ABB yield the same behavior in the Ultimatum Minigame and similar game forms. Say that a game form is a **leader-followers game** if there is only one active player in the first stage, who does not move in stage

¹³Recall that $I(h)$ is the set of active players at h , possibly including chance (see Table 1). For example, $I(\emptyset) = \{c\}$ in the game form of Figure C.

two: $I(\emptyset) = \{j\}$ and $I(\emptyset) \cap I(a^1) = \emptyset$ for some $j \in I$ and every a^1 . Let us write u_{i,θ_i} to make the dependence of u_i on parameter θ_i explicit; then (7) implies:

Remark 1 *In leader-followers games, SA and ABB coincide, that is, $u_{i,\theta_i}^{SA} = u_{i,\theta_i}^{ABB}$ for all θ_i .*

Next, we contrast two specific functional forms for ABB.

Could-have-been blame When frustrated after action profile a^1 , player i considers, for each j , what he would have obtained at most, in expectation, had j chosen differently:

$$\max_{a'_j \in A_j(\emptyset)} \mathbb{E} [\pi_i | (a^1_{-j}, a'_j); \alpha_i].$$

If this could-have-been payoff is more than what i currently expects (that is, $\mathbb{E}[\pi_i | a^1; \alpha_i]$), then i blames j , up to i 's frustration (so (5) holds):

$$B_{ij}(a^1; \alpha_i) = \min \left\{ \left[\max_{a'_j \in A_j(\emptyset)} \mathbb{E} [\pi_i | (a^1_{-j}, a'_j); \alpha_i] - \mathbb{E}[\pi_i | a^1; \alpha_i] \right]^+, F_i(a^1; \alpha_i) \right\}. \quad (8)$$

Blame function (8) satisfies (7) (cf. Remark 4 below).

Example 5 Consider Penny at $a^1 = (D, L)$ in Figure A. For each $j \in \{a, b\}$, Penny's could-have-been payoff is $2 \geq E[\pi_p; \alpha_p]$, her expected payoff is $E[\pi_p | (D, L); \alpha_p] \leq 1$, and her frustration is $[\mathbb{E}[\pi_p; \alpha_p] - 1]^+$. Therefore

$$B_{pa}((D, L); \alpha_p) = B_{pb}((D, L); \alpha_p) = \min \left\{ [2 - \mathbb{E}[\pi_p | (D, L); \alpha_p]]^+, [\mathbb{E}[\pi_p; \alpha_p] - 1]^+ \right\} = [\mathbb{E}[\pi_p; \alpha_p] - 1]^+,$$

that is, both Ann and Bob are fully blamed by Penny for her frustration at (D, L) . \blacktriangle

Blaming unexpected deviations When frustrated after a^1 , i assesses, for each j , how much he would have obtained had j behaved as expected:

$$\sum_{a'_j \in A_j(\emptyset)} \alpha_{ij}(a'_j) \mathbb{E} [\pi_i | (a^1_{-j}, a'_j); \alpha_i],$$

where $\alpha_{ij}(a'_j)$ is the marginal probability of action a'_j according to i 's belief α_i . With this, the blame formula is

$$\begin{aligned} & B_{ij}(a^1; \alpha_i) = \\ & \min \left\{ \left[\sum_{a'_j \in A_j(\emptyset)} \alpha_{ij}(a'_j) \mathbb{E} [\pi_i | (a^1_{-j}, a'_j); \alpha_i] - \mathbb{E}[\pi_i | a^1; \alpha_i] \right]^+, F_i(a^1; \alpha_i) \right\}. \end{aligned} \quad (9)$$

If j is not active in the first stage, we get

$$B_{ij}(a^1; \alpha_i) = \min \left\{ [\mathbb{E}[\pi_i | a^1; \alpha_i] - \mathbb{E}[\pi_i | a^1; \alpha_i]]^+, F_i(a^1; \alpha_i) \right\} = 0;$$

that is, j cannot have deviated and cannot be blamed. If, instead, j is the only active player in the first stage, then

$$\sum_{a'_j \in A_j(\emptyset)} \alpha_{ij}(a'_j) \mathbb{E} [\pi_i | (a^1_{-j}, a'_j); \alpha_i] = \sum_{a' \in A(\emptyset)} \alpha_i(a') \mathbb{E} [\pi_i | a'; \alpha_i] = \mathbb{E} [\pi_i; \alpha_i],$$

and (9) yields

$$B_{ij}(a^1; \alpha_i) = \min \left\{ [\mathbb{E}[\pi_i; \alpha_i] - \mathbb{E}[\pi_i | a^1; \alpha_i]]^+, F_i(a^1; \alpha_i) \right\} = F_i(a^1; \alpha_i).$$

Therefore, like blame function (8), also (9) satisfies (7).

If a^1_j is what i expected j to do in the first stage ($\alpha_{ij}(a^1_j) = 1$) then

$$B_{ij}(a^1; \alpha_i) = \min \left\{ [\mathbb{E}[\pi_i | a^1; \alpha_i] - \mathbb{E}[\pi_i | a^1; \alpha_i]]^+, F_i(a^1; \alpha_i) \right\} = 0.$$

In other words, j did not deviate from what i expected and j is not blamed by i . This is different from “could-have-been” blame (8).

Example 6 Suppose that, in Figure A, Penny is initially certain of (U, L) , so $\alpha_p(U, L) = 1$ and $E[\pi_p; \alpha_p] = 2$. Upon observing (D, L) her frustration is

$F_p((D, L); \alpha_p) = [E[\pi_p; \alpha_p] - 1]^+ = 1$. Using Equation (9), at $a^1 = (D, L)$, Penny fully blames Ann, who deviated from U to D . Using that

$$\sum_{a'_a \in A_a(\emptyset)} \alpha_{pa}(a'_a) \mathbb{E} [\pi_p | (a^1_{-a}, a'_a); \alpha_p] = \pi_p(U, L) = 2$$

we get that Penny's blame of Ann equals Penny's frustration

$$B_{pa}((D, L); \alpha_p) = \min \left\{ [2 - \mathbb{E}[\pi_p | a^1; \alpha_p]]^+, 1 \right\} = 1.$$

On the other hand, Penny does not blame Bob, who played L as expected. To verify this, note that when frustrated after (D, L) Penny assesses how much she would have obtained had Bob behaved as expected:

$$\sum_{a'_b \in A_b(\emptyset)} \alpha_{pb}(a'_b) \mathbb{E} [\pi_p | (a^1_{-b}, a'_b); \alpha_p] = \mathbb{E}[\pi_p | (D, L); \alpha_p]$$

and

$$B_{pb}((D, L); \alpha_p) = \min \left\{ [\mathbb{E}[\pi_p | (D, L); \alpha_p] - \mathbb{E}[\pi_p | (D, L); \alpha_p]]^+, 1 \right\} = 0,$$

in contrast to could-have-been blame (5) under which, as we saw, Penny fully blames Bob (Example 5). \blacktriangle

Blaming unexpected deviations and could-have-been blame both credit the full frustration on the first-mover of a leader-followers game, because they both satisfy (7)(see Remark 1).

4.3 Anger from blaming intentions (ABI)

A player i prone to **anger from blaming intentions (ABI)** asks himself, for each co-player j , whether j intended to give him a low expected payoff. Since such intention depends on j 's first-order beliefs α_j (which include j 's plan, $\alpha_{j,j}$), how much i blames j depends on i 's second-order beliefs β_i , and the decision utility function has the form

$$u_i^{ABI}(h, a_i; \beta_i) = \mathbb{E}[\pi_i | (h, a_i); \alpha_i] - \theta_i \sum_{j \neq i} B_{ij}(h; \beta_i) \mathbb{E}[\pi_j | (h, a_i); \alpha_i],$$

where α_i is derived from β_i .

The maximum payoff that j , initially, can expect to give to i is

$$\max_{a_j^1 \in A_j(\emptyset)} \sum_{a_{-j}^1 \in A_{-j}(\emptyset)} \alpha_{j,-j}(a_{-j}^1) \mathbb{E} [\pi_i | (a_j^1, a_{-j}^1); \alpha_j].$$

Note that

$$\begin{aligned} & \max_{a_j^1 \in A_j(\emptyset)} \sum_{a_{-j}^1 \in A_{-j}(\emptyset)} \alpha_{j,-j}(a_{-j}^1) \mathbb{E} [\pi_i | (a_j^1, a_{-j}^1); \alpha_j] \\ & \geq \sum_{a^1 \in A(\emptyset)} \alpha_j(a^1) \mathbb{E} [\pi_i | a^1; \alpha_j] = \mathbb{E} [\pi_i | \alpha_j], \end{aligned}$$

where the inequality holds by definition, and the equality is implied by the chain rule (2). Note also that $\alpha_j(\cdot | a^1)$ is kept fixed under the maximization; we focus on what j considers he could achieve at the root, taking the view that he cannot control a_j^2 but predicts how he will choose in stage 2. We assume that i 's blame on j at a^1 equals i 's expectation, given second-order belief β_i and conditional on a^1 , of the difference between the maximum payoff that j can expect to give to i and what j plans/expects to give to i , capped by i 's frustration:

$$B_{ij}(a^1; \beta_i) = \tag{10}$$

$$\min \left\{ \mathbb{E} \left[\max_{a_j^1} \sum_{a_{-j}^1} \alpha_{j,-j}(a_{-j}^1) \mathbb{E} [\pi_i | (a_j^1, a_{-j}^1); \alpha_j] - \mathbb{E} [\pi_i; \alpha_j] \middle| a^1; \beta_i \right], F_i(a^1; \alpha_i) \right\},$$

which is non-negative as per the previously highlighted inequality. Now, i 's decision utility after the first-stage history $h = a^1$ is

$$u_i^{ABI}(h, a_i; \beta_i) = \mathbb{E} [\pi_i | (h, a_i); \alpha_i] - \theta_i \sum_{j \neq i} B_{ij}(a^1; \beta_i) \mathbb{E} [\pi_j | (a^1, a_i); \alpha_i],$$

where, in both equations, α_i is derived from β_i .

Example 7 Consider the Ultimatum Minigame form of Figure B. The maximum payoff Ann can give to Bob is 2, independently of α_a . Suppose that Bob, upon observing the greedy offer g , is certain that $\alpha_a(g) = p$, that is, $\beta_b(\alpha_a(g) = p | g) = 1$, with $p < 1$. Also, Bob is certain after g that Ann expected him to accept the greedy offer with probability q , that is, $\beta_b(\alpha_a(y | g) = q | g) = 1$. Finally, suppose Bob initially expected to get the fair

offer ($\alpha_b(f) = 1$), so that his frustration after g is $F_b(a^1; \alpha_b) = 2 - 1 = 1$. Then the extent of Bob’s blame on Ann’s intentions is

$$B_{ba}(g; \beta_b) = \min \{2 - [2(1 - p) + qp], 1\} = \min \{p(2 - q), 1\}.$$

If p is low enough, or q high enough, Bob does not blame all his frustration on Ann. He gives her some credit for the initial intention to make the fair offer with probability $1 - p > 0$, and the degree of credit depends on q . ▲

5 Equilibrium analysis

In this paper we depart from traditional game-theoretic analysis, modeling the role of anger by means of belief-dependent decision-utility functions. With this, our equilibrium analysis is otherwise quite traditional in the following sense: We interpret an equilibrium as a profile of strategies and beliefs representing a “commonly understood” way to play the game by rational (utility maximizing) agents. This is a choice of focus rather than a full endorsement of traditional equilibrium analysis. We want to analyze how equilibrium outcomes change when we take frustration and anger into account.¹⁴

We consider two notions of equilibrium. The first one is the sequential equilibrium (SE) concept of Battigalli & Dufwenberg (2009),¹⁵ extending Kreps & Wilson’s (1982) classic solution to psychological games. In a complete information framework like the one we adopt here for simplicity,¹⁶ SE requires that each player i is certain and never changes his mind about the

¹⁴As stressed by Battigalli & Dufwenberg (2009), with belief-dependent preferences the need to explore and use alternative solution concepts like rationalizability and self-confirming equilibrium is even stronger than with standard preferences. Battigalli & Dufwenberg (2009, Section 5) analyze rationalizability and psychological forward-induction reasoning. Battigalli, Charness & Dufwenberg (2013) apply a notion of incomplete-information rationalizability to show that observed patterns of deceptions can be explained by guilt aversion. Self-confirming equilibrium, instead, has not yet been used in the analysis of psychological games.

¹⁵We consider the version for preferences with own-plan dependence and “local” psychological utility functions (see Battigalli & Dufwenberg 2009, Section 6).

¹⁶Recall that complete information means that the rules of the game and players’ (psychological) preferences are common knowledge. For an illustration of incomplete-information equilibrium analysis of psychological games see, *e.g.*, Attanasi, Battigalli & Manzoni (2015).

true beliefs and plans, hence intentions, of his co-players. We find this feature questionable; therefore, we also explore a generalization – “polymorphic sequential equilibrium” (PSE) – that allows for meaningful updating about others’ intentions.

Battigalli & Dufwenberg’s (2009) SE concept gives equilibrium conditions for infinite hierarchies of conditional probability systems. In our particular application, utility functions only depend on first- or second-order beliefs, so we define SEs for assessments comprising beliefs up to only the second order. Since, technically, first-order beliefs are features of second-order beliefs (see 2.2), we provide definitions that depend only on second-order beliefs, which gives SEs for games where psychological utility functions depend only of first-order beliefs as a special case. Finally, although we so far restrict our analysis of frustration and anger to two-stage game forms, our abstract definitions of equilibrium for games with belief-dependent preferences (and the associated existence theorem) apply to all multistage game forms.

5.1 Sequential equilibrium (SE)

Fix a game form and decision-utility functions $u_i(h, \cdot; \cdot) : A_i(h) \times \Delta_i^2 \rightarrow \mathbb{R}$ ($i \in I, h \in H$). This gives a **psychological game** in the sense of Battigalli & Dufwenberg (2009).¹⁷ An **assessment** is a profile of behavioral strategies and beliefs $(\sigma_i, \beta_i)_{i \in I} \in \times_{i \in I} \Sigma_i \times \Delta_i^2$ such that $\Sigma_i = \times_{h \in H} \Delta(A_i(h))$ and for each $i \in I$, σ_i is the plan $\alpha_{i,i}$ entailed by CPS β_i :

$$\sigma_i(a_i|h) = \alpha_{i,i}(a_i|h) = \beta_i(Z(h, a_i) \times \Delta_{-i}^1|h) \quad (11)$$

for all $i \in I, h \in H, a_i \in A_i(h)$. Eq. (11) implies that the behavioral strategies contained in an assessment are implicitly determined by players’ beliefs about paths; therefore, they could be dispensed with. We follow Battigalli & Dufwenberg (2009) and make behavioral strategies explicit in assessments only to facilitate understanding and comparisons with the equilibrium refinements literature.

Definition 1 *An assessment $(\sigma_i, \beta_i)_{i \in I}$ is **consistent** if, for all $i \in I, h \in H$, and $a = (a_j)_{j \in I(h)} \in A(h)$,*

¹⁷See the extensions discussed in Section 6. Note that we are not modeling beliefs about co-players’ sensitivity to anger (nor beliefs about such beliefs, etc.), because we are assuming common knowledge of psychological preferences.

(a) $\alpha_i(a|h) = \times_{j \in I(h)} \sigma_j(a_j|h)$,
(b) $\text{marg}_{\Delta_{-i}^1} \beta_i(\cdot|h) = \delta_{\alpha_{-i}}$,
where α_j is derived from β_j for each $j \in I$, and $\delta_{\alpha_{-i}}$ is the Dirac probability measure that assigns probability one to the singleton $\{\alpha_{-i}\} \subseteq \Delta_{-i}^1$.

Condition (a) requires that players' beliefs about actions satisfy independence across co-players (on top of own-action independence), and each i expects each j to behave as specified by j 's plan $\sigma_j = \alpha_{j,j}$. This implies that all players have the same first-order beliefs. Condition (b) requires that players' beliefs about co-players' first-order beliefs (hence their plans) are correct and never change, on or off the path. This implies that all players, essentially, have the same second-order beliefs (considering that they are introspective and therefore know their own first-order beliefs). For later use, we record two facts about consistency:

Lemma 1 *For each behavioral strategies profile $\sigma = (\sigma_i)_{i \in I}$ there is a unique profile of second-order beliefs $\beta^\sigma = (\beta_i^\sigma)_{i \in I}$ such that (σ, β^σ) is a consistent assessment. The map $\sigma \mapsto \beta^\sigma$ is continuous.*

Proof Write $\mathbb{P}^\sigma(h'|h)$ for the probability of reaching h' from h , e.g.,

$$\mathbb{P}^\sigma(a^1, a^2|\emptyset) = \left(\prod_{j \in I(\emptyset)} \sigma_j(a_j^1|\emptyset) \right) \left(\prod_{j \in I(a^1)} \sigma_j(a_j^2|a^1) \right).$$

With this, first define α_i^σ as $\alpha_i^\sigma(z|h) = \mathbb{P}^\sigma(z|h)$ for all $i \in I$, $h \in H$, and $z \in Z$. Next define β^σ as $\beta_i^\sigma(\cdot|h) = \alpha_i^\sigma(\cdot|h) \times \delta_{\alpha_{-i}^\sigma}$ for all $i \in I$, $h \in H$. It can be checked that (1) $\beta^\sigma \in \Delta_i^2$ for each $i \in I$, (2) (σ, β^σ) is a consistent assessment, and (3) if $\beta \neq \beta^\sigma$, then either (a) or (b) of Definition 1 is violated. It is also apparent from the construction that the map $\sigma \mapsto \beta^\sigma$ is continuous, because $\sigma \mapsto \alpha^\sigma$ is obviously continuous, and the Dirac-measure map $\alpha_{-i} \mapsto \delta_{\alpha_{-i}}$ is continuous. ■

Lemma 2 *The set of consistent assessments is compact.*

Proof The set of consistent assessments is contained in the compact metrizable space $\times_{i \in I} (\Sigma_i \times \Delta_i^2)$. Therefore, it is enough to show that it is closed. Let $(\sigma^n, \beta^n)_{n \in \mathbb{N}}$ be a converging sequence of consistent assessments

with limit $(\sigma^\infty, \beta^\infty)$. For each $i \in I$, let α_i^n be the first-order belief derived from β_i^n ($n \in \mathbb{N} \cup \{\infty\}$), that is,

$$\alpha_i^n(Y|h) = \beta_i^n(Y \times \Delta_{-i}^1|h)$$

for all $h \in H$ and $Y \subseteq Z(h)$. By consistency, for all $n \in \mathbb{N}$, $i \in I$, $h \in H$, $a \in A(h)$, and $E_{-i} \subseteq \Delta_{-i}$ it holds that

- (a.n) $\alpha_i^n(a|h) = \beta_i^n(Z(h, a) \times \Delta_{-i}^1|h) = \prod_{j \in I(h)} \sigma_j^n(a_j|h)$,
- (b.n) $\text{marg}_{\Delta_{-i}^1} \beta_i^n(\cdot|h) = \delta_{\alpha_{-i}^n}$, where each α_{-i}^n is determined by σ^n as per (a.n).

Then, of course, $\alpha_i^\infty(a|h) = \beta_i^\infty(Z(h, a) \times \Delta_{-i}^1|h) = \prod_{j \in I(h)} \sigma_j^\infty(a_j|h)$ for all $i \in I$, $h \in H$, $a \in A(h)$. Furthermore, $\text{marg}_{\Delta_{-i}^1} \beta_i^\infty(\cdot|h) = \delta_{\alpha_{-i}^\infty}$ for all $i \in I$ and $h \in H$, because $\alpha_{-i}^n \rightarrow \alpha_{-i}^\infty$ and the marginalization and Dirac maps $\beta_i \mapsto \text{marg}_{\Delta_{-i}^1} \beta_i$ and $\alpha_{-i} \mapsto \delta_{\alpha_{-i}}$ are continuous. ■

Definition 2 *An assessment $(\sigma_i, \beta_i)_{i \in I}$ is a **sequential equilibrium (SE)** if it is consistent and satisfies the following sequential rationality condition: for all $h \in H$ and $i \in I(h)$*

$$\text{Supp} \sigma_i(\cdot|h) \subseteq \arg \max_{a_i \in A_i(h)} u_i(h, a_i; \beta_i).$$

It can be checked that this definition of SE is equivalent to the traditional one when players have standard preferences, that is, when there is a profile of utility functions $(v_i : Z \rightarrow \mathbb{R})_{i \in I}$ such that $u_i(h, a_i; \beta_i) = \mathbb{E}[v_i|(h, a_i); \alpha_i]$.¹⁸ A special case of this is the **material-payoff game**, where $v_i = \pi_i$ for each $i \in I$.

Theorem 3 *If $u_i(h, a_i; \cdot)$ is continuous for all $i \in I$, $h \in H$ and $a_i \in A_i(h)$, then there is at least one SE.*

Battigalli & Dufwenberg (2009) prove a version of this existence result where first-order beliefs are modeled as CPSs over pure strategies profiles

¹⁸According to the standard definition of SE, sequential rationality is given by global maximization over (continuation) strategies at each $h \in H$. By the One-Shot-Deviation principle, this is equivalent to “local” maximization over actions at each $h \in H$.

rather than paths. But their “trembling-hand” technique can be used here with straightforward adaptations. We omit the details.¹⁹

What we said so far about equilibrium does not assume specific functional forms. From now on, we focus on u_i^{SA} , u_i^{ABB} , and u_i^{ABI} . Since frustration and blame are continuous in beliefs, decision-utility is also continuous, and we obtain existence in all cases of interest:

Corollary 1 *Every game with SA, ABB, or ABI has at least one SE.*

Remark 2 *Let $(\sigma_i, \beta_i)_{i \in I}$ be a SE assessment of a game with SA, ABB, or ABI; if a history $h \in H$ has probability one under profile $(\sigma_i)_{i \in I}$, then*

$$F_i(h'; \alpha_i) = 0, \text{ Supp}\sigma_i(\cdot|h') \subseteq \arg \max_{a'_i \in A_i(h')} \mathbb{E}[\pi_i|h'; \alpha_i]$$

for all $h' \preceq h$ and $i \in I(h')$, where α_i is derived from β_i . Therefore, a SE strategy profile of a game with SA, ABB, or ABI with randomization (if any) only in the last stage is also a Nash equilibrium of the agent form of the corresponding material-payoff game.

Proof Fix $i \in I$ arbitrarily. First-order belief α_i is derived from β_i and, by consistency, gives the behavioral strategies profile σ . Therefore, by assumption each $h' \preceq h$ has probability one under α_i , which implies that $\mathbb{E}[\pi_i|h'; \alpha_i] = \mathbb{E}[\pi_i; \alpha_i]$, hence $F_i(h'; \alpha_i) = 0$. Since blame is capped by frustration, $u_i(h', a'_i; \beta_i) = \mathbb{E}[\pi_i|h'; \alpha_i]$. Therefore, sequential rationality of the equilibrium assessment implies that $\text{Supp}\sigma_i(\cdot|h') \subseteq \arg \max_{a'_i \in A_i(h')} \mathbb{E}[\pi_i|h'; \alpha_i]$ if $i \in I(h')$. If there is randomization (if any) only in the last state, then players maximize (locally) their expected material payoff on the equilibrium path. Hence, the second claim follows by inspection of the definitions of agent form (of the material-payoff game) and Nash equilibrium. ■

To illustrate, in the Ultimatum minigame (f, n) can be a SE under ABB, and is also a Nash equilibrium of the agent form with material-payoff utilities. Essentially, with (counterfactual) anger in the picture, n becomes a credible threat. Corollary 1 and Remark 2 also hold for the multistage extension of Section 6.

We say that two assessments are **realization-equivalent** if the corresponding strategy profiles yield the same probability distribution over terminal histories.

¹⁹A similar technique is used in the first part of the proof of Proposition 1.

Proposition 1 *In every perfect-information (two-stage) game form without chance moves and a unique SE of the material-payoff game, this equilibrium is realization-equivalent to a SE of the psychological game with ABI, ABB, or - with only two players, SA.*

Proof Let $(\bar{\sigma}, \bar{\beta}) = (\bar{\sigma}_i, \bar{\beta}_i)_{i \in I}$ be the SE of the material payoff game, which is in pure strategies by the perfect information assumption. Fix decision-utility functions $u_i(h, a_i; \cdot)$ of the ABI, or ABB kind, and a sequence of real numbers $(\varepsilon_n)_{n \in \mathbb{N}}$, with $\varepsilon_n \rightarrow 0$ and $0 < \varepsilon_n < \frac{1}{\max_{i \in I, h \in H} |A_i(h)|}$ for all $n \in \mathbb{N}$. Consider the constrained psychological game with such decision-utility functions where players can choose mixed actions in the following sets:

$$\Sigma_i^n(h) = \{\sigma_i(\cdot|h) \in \Delta(A_i(h)) : \|\sigma_i(\cdot|h) - \bar{\sigma}_i(\cdot|h)\| \leq \varepsilon_n\}$$

if h is on the $\bar{\sigma}$ -path, and

$$\Sigma_i^n(h) = \{\sigma_i(\cdot|h) \in \Delta(A_i(h)) : \forall a_i \in A_i(h), \sigma_i(a_i|h) \geq \varepsilon_n\}$$

if h is off the $\bar{\sigma}$ -path. By construction, these sets are non-empty, convex, and compact. Since the decision-utility functions are continuous in beliefs, and the consistent assessment map $\sigma \mapsto \beta^\sigma$ is continuous (Lemma 1), the correspondence

$$\sigma \mapsto \times_{h \in H} \times_{i \in I(h)} \arg \max_{\sigma_i(\cdot|h) \in \Sigma_i^n(h)} \sum_{a_i \in A_i(h)} \sigma'_i(a_i|h) u_i(h, a_i; \beta_i^\sigma)$$

is upper-hemicontinuous, non-empty, convex, and compact valued; therefore (by Kakutani's theorem), it has a fixed point σ^n . By Lemma 2, the sequence of consistent assessments $(\sigma^n, \beta^{\sigma^n})_{n=1}^\infty$ has a limit point (σ^*, β^*) , which is consistent too. By construction, $\bar{\sigma}(\cdot|h) = \sigma^*(\cdot|h)$ for h on the $\bar{\sigma}$ -path, therefore $(\bar{\sigma}, \bar{\beta})$ and (σ^*, β^*) are realization-equivalent. We let $\bar{\alpha}_i$ (respectively, α_i^*) denote the first-order beliefs of i implied by $(\bar{\sigma}, \bar{\beta})$ (respectively, by (σ^*, β^*)).

We claim that the consistent assessment (σ^*, β^*) is a SE of the psychological game with decision-utility functions $u_i(h, a_i; \cdot)$. We must show that (σ^*, β^*) satisfies sequential rationality. If h is off the $\bar{\sigma}$ -path, sequential rationality is satisfied by construction. Since $\bar{\sigma}$ is deterministic and there are no chance moves, if h is on the $\bar{\sigma}$ -path (that is on the σ^* -path) it must have unconditional probability one according to each player's beliefs and there cannot be any frustration; hence, $u_i(h, a_i; \beta_i^*) = \mathbb{E}[\pi_i|h, a_i; \alpha_i^*]$ ($i \in I$) where

α_i^* is determined by σ^* . If, furthermore, it is the second stage ($h = \bar{a}^1$, with $\bar{\sigma}(\bar{a}^1|\emptyset) = 1$), then – by construction – $\mathbb{E}[\pi_i|h, a_i; \alpha_i^*] = \mathbb{E}[\pi_i|h, a_i; \bar{\alpha}_i]$, where $\bar{\alpha}_i$ is determined by $\bar{\sigma}$. Since $\bar{\sigma}$ is a SE of the material-payoff game, sequential rationality is satisfied at h . Finally, we claim that (σ^*, β^*) satisfies sequential rationality also at the root $h = \emptyset$. Let $\iota(h)$ denote the active player at h . Since $\iota(\emptyset)$ cannot be frustrated at \emptyset , we must show that action \bar{a}^1 with $\bar{\sigma}(\bar{a}^1|\emptyset) = 1$ maximizes his expected material payoff given belief $\alpha_{\iota(\emptyset)}$. According to ABB and ABI, player $\iota(a^1)$ can only blame the first mover $\iota(\emptyset)$ and possibly hurt him, if he is frustrated. Therefore, in assessment (σ^*, β^*) at node a^1 , either $\iota(a^1)$ plans to choose his (unique) payoff maximizing action, or he blames $\iota(\emptyset)$ strongly enough to give up some material payoff in order to bring down the payoff of $\iota(\emptyset)$. Hence, $\mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \alpha_{\iota(a^1)}^*] \leq \mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \bar{\alpha}_{\iota(a^1)}]$ (anger). By consistency of (σ^*, β^*) and $(\bar{\sigma}, \bar{\beta})$, $\alpha_{\iota(a^1)}^* = \alpha_{\iota(\emptyset)}^*$ and $\bar{\alpha}_{\iota(a^1)} = \bar{\alpha}_{\iota(\emptyset)}$ (cons.). Since (σ^*, β^*) is realization-equivalent to $(\bar{\sigma}, \bar{\beta})$ (r.e.), which is the material-payoff equilibrium (m.eq.), for each $a^1 \in A(\emptyset)$,

$$\begin{aligned} \mathbb{E}[\pi_{\iota(\emptyset)}|\bar{a}^1; \alpha_{\iota(\emptyset)}^*] &\stackrel{\text{(r.e.)}}{=} \mathbb{E}[\pi_{\iota(\emptyset)}|\bar{a}^1; \bar{\alpha}_{\iota(\emptyset)}] \stackrel{\text{(m.eq.)}}{\geq} \\ \mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \bar{\alpha}_{\iota(\emptyset)}] &\stackrel{\text{(cons.)}}{=} \mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \bar{\alpha}_{\iota(a^1)}] \stackrel{\text{(anger)}}{\geq} \\ \mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \alpha_{\iota(a^1)}^*] &\stackrel{\text{(cons.)}}{=} \mathbb{E}[\pi_{\iota(\emptyset)}|a^1; \alpha_{\iota(\emptyset)}^*]. \end{aligned}$$

This completes the proof that (σ^*, β^*) is a SE if the decision-utility functions are of the ABB or ABI kind. If there are only two players, then we have a leader-follower game and SA is equivalent to ABB (Remark 1). Therefore, (σ^*, β^*) is a SE in this case too. ■

We note that the assumption of a unique material-payoff SE holds generically in game forms with perfect information. It is quite easy to show by example that without perfect information, or with chance moves, a material-payoff SE need not be a SE with frustration and anger. The same holds for some multistage game forms (see the analysis of Section 6). Disregarding chance moves, randomization, and ties, the common feature of material-payoff equilibria that are not realization-equivalent to equilibria with frustration and anger is the following (see also Example 9 below): An off-path frustrated player j wants to hurt co-player k which implies rewarding a preceding on-path player i ; this makes it impossible to satisfy both i 's incentive to not deviate and j 's incentive to punish.

We close this section with three examples, which combine to illustrate how our SE concept works (including a weakness) and that the notions of SA, ABB (both versions), and ABI may alter material incentives, and may produce different predictions.

Example 8 Consider Figure C (“Hammering one’s thumb”). With u_a^{ABB} (either version), or u_a^{ABI} , Andy will not blame Bob so his SE-choice is the material-payoff equilibrium, N . But with u_a^{SA} Andy may take it out on Bob (*i.e.*, choose T). Recall that $F_a(B; \alpha_a) = 2(1 - \varepsilon) + \varepsilon\alpha_a(N|B) - 1$, so the more likely Andy believes it to be that he will take it out on Bob, the less he expects initially and the less frustrated he is when B happens. Yet, in SE, the more prone to get angry he is (as measured by θ_a) the more likely that he will take it out on Bob: Andy’s utility from N and T is

$$\begin{aligned} u_a^{SA}(B, N; \alpha_i) &= 1 - \theta_a[2(1 - \varepsilon) + \varepsilon\alpha_a(N|B) - 1] \cdot 1, \\ u_a^{SA}(B, T; \alpha_i) &= 0 - \theta_a[2(1 - \varepsilon) + \varepsilon\alpha_a(N|B) - 1] \cdot 0 = 0. \end{aligned}$$

Sequential rationality of SE implies that one possibility is $\alpha_a(N|B) = 1$ and $u_a^{SA}(B, N; \alpha_i) \geq u_a^{SA}(B, T; \alpha_i)$, implying $\theta_a \leq \frac{1}{1-\varepsilon}$. Another possibility is $\alpha_a(N|B) = 0$ and $u_a^{SA}(B, N; \alpha_i) \leq u_a^{SA}(B, T; \alpha_i)$, implying $\theta_a \leq \frac{1}{1-2\varepsilon}$. If $\theta_a \in (\frac{1}{1-\varepsilon}, \frac{1}{1-2\varepsilon})$, we can solve for a SE where $u_a^{SA}(B, N; \alpha_i) = u_a^{SA}(B, T; \alpha_i)$ and $\alpha_a(N|B) = \frac{1}{\varepsilon\theta_a} - \frac{1-2\varepsilon}{\varepsilon} \in (0, 1)$. \blacktriangle

The case where $\theta_a \in (\frac{1}{1-\varepsilon}, \frac{1}{1-2\varepsilon})$ illustrates how we cannot take for granted that a SE exists where players use deterministic plans (a point relevant also for u_i^{ABB} or u_i^{ABI} in other games). Here this happens in a game form with a single active player, highlighting that we deal with a psychological game, as this could not be the case in a standard game.

Example 9 Consider Figure A (“Asymmetric Punishment”). Can the material-payoff equilibrium outcome (U, L) be part of a SE with frustration and anger? The answer is yes under ABI and the blaming-unexpected-deviations version of ABB. To see this note that Ann and Bob act as-if selfish (as they are not frustrated). Hence they would deviate if they could gain material payoff. In the SE, they would expect 5 if not deviating, making Ann the sole deviation candidate (she’d get $6 > 5$ were Penny to choose P ; for Bob, 5 is the best he could hope for). Ann deviating can be dismissed though, since if (D, L) were reached Penny would not blame Bob (the only co-player she can punish)

under either relevant blame function, and so she would choose N (regardless of θ_p). Under SA and the could-have-been version of ABB, however, it may be impossible to sustain a SE with (U, L) ; at (D, L) Penny would blame each of Ann and Bob (as explained earlier). By choosing P she hurts Bob more than she helps Ann and would do so if

$$u_p^{ABB}((D, L), P; \alpha_p) > u_p^{ABB}((D, L), N; \alpha_p)$$

$$\iff$$

$$0 - 6\theta_p B_{pa}((D, L); \alpha_p) > 1 - 8\theta_p B_{pa}((D, L); \alpha_p).$$

The rhs of the last inequality uses $B_{pb}((D, L); \alpha_p) = B_{pa}((D, L); \alpha_p)$. Since $B_{pa}((D, L); \alpha_p) = F_p((D, L); \alpha_i) = 1 > 0$, Penny would choose P if $-6\theta_p > 1 - 8\theta_p \iff \theta_p > 1/2$, so Ann would want to deviate and choose D . \blacktriangle

Example 10 Consider Figure B (“Ultimatum Minigame”). By Proposition 1, every utility function discussed admits the unique material-payoff equilibrium (g, y) as a SE, regardless of anger sensitivity. To check this directly, just note that, if Bob expects g , he cannot be frustrated, so – when asked to play – he maximizes his material payoff. Under SA and ABB (both versions), (f, n) qualifies as another SE if $\theta_b \geq 1/3$; following g , Bob would be frustrated and choose n , so Ann chooses f . Under ABI (f, n) cannot be an SE. To verify, assume it were, so $\alpha_a(f) = 1$. Since the SE concept does not allow for players revising beliefs about beliefs, we get $\beta_b(\alpha_a(f) = 1|g) = 1$ and $B_{ba}(g; \beta_b) = 0$; Bob maintains his belief that Ann planned to choose f , hence she intended to maximize Bob’s payoff. Hence, Bob would choose y , contradicting that (f, n) is a SE. Next, note that (g, n) is not a SE under any concept: Given SE beliefs Bob would not be frustrated and so he would choose y . With ABI, the only way to observe rejected offers with positive probability in a SE is with non-deterministic plans. To find such a SE, note that we need $\alpha_a(g) \in (0, 1)$; if $\alpha_a(g) = 0$ Bob would not be reached and if $\alpha_a(g) = 1$ he would not be frustrated, and hence, he would choose y . Since Ann uses a non-degenerate plan she must be indifferent, so $\alpha_b(y) = 2/3$, implying that Bob is indifferent too. In SE, Bob’s frustration is $[2(1 - \alpha_a(g)) + \frac{2}{3}\alpha_a(g) - 1]^+ = [1 - \frac{4}{3}\alpha_a(g)]^+$, which equals his blame of

Ann under SA and ABB. Hence we get the indifference condition

$$\begin{aligned} 1 - \theta_b \left[1 - \frac{4}{3}\alpha_a(g) \right]^+ \cdot 3 &= 0 - \theta_b \left[1 - \frac{4}{3}\alpha_a(g) \right]^+ \cdot 0 \\ &\iff \\ \alpha_a(g) &= \frac{3}{4} - \frac{1}{4\theta_b}, \end{aligned}$$

where $\theta_b \geq 1/3$. The more prone to anger Bob is the more likely he is to get the low offer, so Bob's initial expectations, and hence his frustration and blame, is kept low. Under ABI we get another indifference condition:

$$\begin{aligned} 1 - \theta_b B_{ba}(g; \beta_b) \cdot 3 &= 0 - \theta_b B_{ba}(g; \beta_b) \cdot 0 \\ &\iff \\ 1 - \theta_b \min \left\{ 1 - \frac{4}{3}\alpha_a(g), \frac{4}{3}\alpha_a(g) \right\} \cdot 3 &= 0. \end{aligned}$$

The left term in braces is Bob's frustration while

$$\frac{4}{3}\alpha_a(g) = 2 - \left[2(1 - \alpha_a(g)) + \frac{2}{3}\alpha_a(g) \right]$$

is the difference between the maximum payoff Ann could plan for Bob and her actually planned one. The first term is lower if $\alpha_a(g) \geq 3/8$; so, if we can solve the equation for such a number, we duplicate the SA/ABB-solution; again, this is doable if $\theta_b > 1/3$. If $\theta_b \geq 2/3$, with ABI, there is second non-degenerate equilibrium plan with $\alpha_a(g) \in (0, \frac{3}{8})$ such that $\alpha_a(g) = 1/4\theta_b$; to see this, solve the ABI indifference condition assuming that $\frac{4}{3}\alpha_a(g) \leq 1 - \frac{4}{3}\alpha_a(g)$. This SE exhibits starkly different comparative statics; the more prone to anger Bob is, the less likely he is to get a low offer and the less he blames Ann following g in light of her intention to choose f with higher probability. \blacktriangle

In the last example we explained why with ABI (f, n) cannot be an SE. We find the interpretation unappealing. If Bob initially expects Ann to choose f , and she doesn't, so that Bob is frustrated, then he would rate her choice a mistake and not blame her! It may seem more plausible for Bob not to be so gullible, and instead revise his beliefs of Ann's intentions. The SE concept rules that out. Because of this, and because it makes sense regardless, we next define an alternative concept that to a degree overcomes the issue.

5.2 Polymorphic sequential equilibrium (PSE)

Suppose a game is played by agents drawn at random and independently from large populations, one for each player role $i \in I$. Different agents in the same population i have the same belief-dependent preferences,²⁰ but they may have different plans, hence different beliefs about paths, even if their beliefs agree about the behavior and beliefs of co-players $-i$. In this case, we say that the population is “polymorphic.” Once an agent playing in role i observes some moves of co-players, he makes inferences about the intentions of the agents playing in the co-players’ roles.

Let λ_i be a finite support distribution over $\Sigma_i \times \Delta_i^2$, with $\text{Supp}\lambda_i = \{(\sigma_{t_i^1}, \beta_{t_i^1}), (\sigma_{t_i^2}, \beta_{t_i^2}), \dots\}$. We interpret λ_i as a statistical distribution of plans and beliefs of agents playing in role i and, for every $(\sigma_{t_i}, \beta_{t_i}) \in \text{Supp}\lambda_i$, we let λ_{t_i} denote the fraction of agents in population i with plan and beliefs $(\sigma_{t_i}, \beta_{t_i})$.²¹ We refer to such index t_i as a “type” of i .²² Also, we denote by

$$T_i(\lambda_i) = \{t_i : (\sigma_{t_i}, \beta_{t_i}) \in \text{Supp}\lambda_i\}$$

the set of possible types of i in distribution λ_i . Also, we write $T_{-i}(\lambda_{-i}) = \times_{j \neq i} T_j(\lambda_j)$ for the set of profiles of co-players’ types.

Let us take the perspective of an agent of type t_i who knows that the distribution over co-players’ types is $\lambda_{-i} = \prod_{j \neq i} \lambda_j$ and believes that the behavior of each t_j is indeed described t_j ’s plan σ_{t_j} (in principle, t_i may otherwise believe that t_j behaves differently from his plan). Then it is possible to derive the conditional probability of a type profile t_{-i} given history h . Given that beliefs satisfy independence across players (everybody knows that there is independent random matching), the distribution is independent of t_i and can be factorized. In the current two-stage setting we have:

$$\lambda_{-i}(t_{-i}|\emptyset) = \prod_{j \neq i} \lambda_j(t_j|\emptyset) = \prod_{j \neq i} \lambda_{t_j},$$

²⁰Recall that we are not modelling incomplete information.

²¹The marginal of λ_i on Σ_i is a behavior strategy mixture (see Selten 1975).

²²They are “types” in the sense of epistemic game theory (*e.g.* Battigalli, Di Tillio & Samet 2013).

and

$$\begin{aligned}
\lambda_{-i}(t_{-i}|a^1) &= \frac{\prod_{j \neq i} \sigma_{t_j}(a_j^1) \lambda_{t_j}}{\sum_{t'_{-i} \in T_{-i}(\lambda_{-i})} \prod_{j \neq i} \sigma_{t'_j}(a_j^1) \lambda_{t'_j}} \\
&= \frac{\prod_{j \neq i} \sigma_{t_j}(a_j^1) \lambda_{t_j}}{\prod_{j \neq i} \sum_{t'_j \in T_j(\lambda_j)} \sigma_{t'_j}(a_j^1) \lambda_{t'_j}} \\
&= \prod_{j \neq i} \frac{\sigma_{t_j}(a_j^1) \lambda_{t_j}}{\sum_{t'_j \in T_j(\lambda_j)} \sigma_{t'_j}(a_j^1) \lambda_{t'_j}}.
\end{aligned}$$

for all t_{-i} and a^1 , provided that $\sum_{t'_j} \sigma_{t'_j}(a_j^1) \lambda_{t'_j} > 0$ for each $j \neq i$. Letting

$$\lambda_j(t_j|a^1) = \frac{\sigma_{t_j}(a_j^1) \lambda_{t_j}}{\sum_{t'_j \in T_j(\lambda_j)} \sigma_{t'_j}(a_j^1) \lambda_{t'_j}},$$

we get

$$\lambda_{-i}(t_{-i}|a^1) = \prod_{j \neq i} \lambda_j(t_j|a^1).$$

We say that λ_j is **fully randomized** if σ_{t_j} is strictly positive for every type $t_j \in T_j(\lambda_j)$. If each λ_j is fully randomized, then, for all $h \in H$, $\lambda_{-i}(\cdot|h)$ is well defined, with $\lambda_{-i}(t_{-i}|h) = \prod_{j \neq i} \lambda_j(t_j|h)$ for all $t_{-i} \in T_{-i}(\lambda_{-i})$.

Definition 3 A *polymorphic assessment* is a profile of finite support probability measures $\lambda = (\lambda_i)_{i \in I} \in \times_{i \in I} \Delta(\Sigma_i \times \Delta_i^2)$ such that, for every $i \in I$ and $t_i \in T_i(\lambda_i)$, σ_{i,t_i} is the behavior strategy obtained from β_{t_i} as per (11). A polymorphic assessment λ is **consistent** if there is a sequence $(\lambda^n)_{n=1}^\infty$ of polymorphic assessments converging to λ such that, for all $j \in I$ and $n \in \mathbb{N}$, λ_j^n is fully randomized, and (a-p) for all $h \in H$, $a \in A(h)$, and $t_i \in T_i(\lambda_i^n)$,

$$\alpha_{t_i, -i}^n(a_{-i}|h) = \prod_{j \neq i} \sum_{t_j \in T_j(\lambda_j^n)} \sigma_{t_j}^n(a_j|h) \lambda_j^n(t_j|h).$$

(b-p) for all $h \in H$ and $t_i \in T_i(\lambda_i^n)$,

$$\text{marg}_{\Delta_{-i}^1} \beta_{t_i}^n(\cdot|h) = \sum_{t_{-i} \in T_{-i}(\lambda_{-i}^n)} \lambda_{-i}^n(t_{-i}|h) \delta_{\alpha_{t_{-i}}^n},$$

where, for all $j \in I$, $t_j \in T_j(\lambda_j^n)$ and $n \in \mathbb{N}$, $\alpha_{t_j}^n$ is the first-order CPS derived from $\beta_{t_j}^n$.

Condition (a-p) extends the independence condition (a) of Definition 1 to the multiple-types setting. Condition (b-p) implies that, conditional on the co-players' types, everyone has correct beliefs about the beliefs of others, including their plans, but uncertainty about co-players' types allows for uncertainty and meaningful updating about such beliefs. Conditions (a-p) and (b-p) imply that different types of the same player share the same beliefs about co-players, but may have different plans. Therefore, Definition 3 is a minimal departure from the notion of consistent assessment, allowing for uncertainty and meaningful updating about the plans, hence intentions, of co-players.

Definition 4 A polymorphic assessment λ is a **polymorphic sequential equilibrium (PSE)** if it is consistent and satisfies the following sequential rationality condition: for all $h \in H$, $i \in I(h)$, and $t_i \in T_i(\lambda_i)$,

$$\text{Supp}\sigma_{t_i}(\cdot|h) \subseteq \arg \max_{a_i \in A_i(h)} u_i(h, a_i; \beta_{t_i}).$$

Remark 3 Every SE is a degenerate (or monomorphic) PSE. Therefore, Theorem 3 implies that, if every decision-utility function $u_i(h, \cdot; \cdot)$ ($i \in I$, $h \in H$) is continuous, then there is at least one PSE. In particular, every game with SA, ABB, or ABI has at least one PSE.

Finally, we demonstrate how the PSE alters predictions in the Ultimatum Minigame and in leader-followers games more generally.

Example 11 Consider again the Ultimatum Minigame in Figure B. If $|\text{Supp}\lambda_i| = 1$ for all i , then our results for the SE analysis still hold as a special case of the more general PSE analysis. Interesting new possibilities arise if $|\text{Supp}\lambda_a| = 2$ though. Recall that, in the SE with SA/ABB utility functions and non-degenerate plans of Example 10, we had $\alpha_a(g) = \frac{3}{4} - \frac{1}{4\theta_b}$ (with $\theta_b > 1/3$) to keep Bob indifferent. Suppose instead there are two types of Ann, a fraction of $\frac{3}{4} - \frac{1}{4\theta_b}$ of them planning to choose g while the others plan for f . There is a corresponding PSE where (naming Ann's types by planned choice) $\text{Supp}\lambda_a = \{(f, \beta_f), (g, \beta_g)\}$, $\alpha_f(y|g) = \alpha_g(y|g) = \alpha_b(y|g) = 2/3$, and this holds for also for ABI, not only SA and ABB. The first-order belief of type f of Ann, α_f , is derived from β_f , etc. Bob initially believes Ann is either an f -type or a g -type, assigning probability $\lambda_g = \frac{3}{4} - \frac{1}{4\theta_b}$ to the latter possibility. Following g he ceases to assign positive probability to being matched

with an f -type, assigning instead probability 1 to the g -type, a meaningful form of updating about Ann's intentions implied by consistency (Def. 3). This inference makes ABI work exactly as ABB (and SA). Bob's frustration is as in Example 10, so equal to his blame of Ann for each blaming function. So again Bob is indifferent between y and n , and sequentially rational if $\alpha_b(y|g) = 2/3$. Condition $\alpha_f(y|g) = \alpha_g(y|g) = 2/3$ implies that both types of Ann are indifferent, hence sequentially rational. Thus, starting with the non-degenerate SE under ABB (and SA) we obtained a PSE, under every blaming function, where Ann's plan is purified.

The observation of the previous example can be generalized to all leader-followers games:²³

Proposition 2 *Consider a leader-followers game and an arbitrary parameter profile $(\theta_i)_{i \in I}$. Every SE of the psychological game with decision-utility functions $(u_{i,\theta_i}^{ABB})_{i \in I}$ [or $(u_{i,\theta_i}^{SA})_{i \in I}$] where the behavioral strategy of the leader has full support corresponds to a PSE of the psychological game with decision-utility functions $(u_{i,\theta_i}^{ABI})_{i \in I}$ and $(u_{i,\theta_i}^{ABB})_{i \in I}$ [or $(u_{i,\theta_i}^{SA})_{i \in I}$] where the leader is purified.*

Proof We denote the leader by $\iota(\emptyset)$. Let $(\sigma_i, \beta_i)_{i \in I}$ be a SE under ABB/SA with parameter profile $(\theta_i)_{i \in I}$, and suppose that $\text{Supp}\sigma_{\iota(\emptyset)}(\cdot|\emptyset) = A_{\iota(\emptyset)}(\emptyset)$. Construct a polymorphic consistent assessment $\bar{\lambda}$ as follows: For each follower i , $T_i(\bar{\lambda}_i) = \{t_i\}$ (a singleton) and $(\bar{\sigma}_{t_i}, \bar{\beta}_{t_i}) = (\sigma_i, \beta_i)$. For the leader $\iota(\emptyset)$, $T_{\iota(\emptyset)}(\bar{\lambda}_{\iota(\emptyset)}) = A_{\iota(\emptyset)}(\emptyset)$, and, for each type $a_{\iota(\emptyset)}$, $\bar{\sigma}_{a_{\iota(\emptyset)}}(a_{\iota(\emptyset)}|\emptyset) = 1$ and $\bar{\alpha}_{a_{\iota(\emptyset)}}(\cdot|a^1) = \sigma_{-i}(\cdot|a^1)$ for all non-terminal a^1 , where $\bar{\alpha}_{a_{\iota(\emptyset)}}$ is the first-order belief derived from $\bar{\beta}_{a_{\iota(\emptyset)}}$. By construction, each type of leader is indifferent, because the leader (who acts as-if selfish) is indifferent in the original assessment $(\sigma_i, \beta_i)_{i \in I}$. As for the followers, they have the same first-order beliefs, hence the same second-stage frustrations as in $(\sigma_i, \beta_i)_{i \in I}$. Under ABB/SA, blame always equals frustration in leader-followers games. As for ABI, Bayes' rule implies that, after observing $a^1 = a_{\iota(\emptyset)}$, each follower becomes certain that the leader indeed planned to choose $a_{\iota(\emptyset)}$ with probability one, and blame equals frustration in this case too. Therefore, the incentive conditions of the followers hold in $\bar{\lambda}$ as in $(\sigma_i, \beta_i)_{i \in I}$ for all kinds of decision utility (ABI, ABB, SA) under the same parameter profile $(\theta_i)_{i \in I}$. ■

²³Recall that, by Remark 1, SA is equivalent to both versions of ABB in such games.

6 Multistage extension

In a multistage game form, a (non-empty) non-terminal history is a sequence of action profiles, $h = (a^1, \dots, a^t)$ where $t \geq 1$. As in the two-stage case, we assume that actions are observable; hence, every non-terminal history is public. Our notation for the multistage setting is essentially the same as before (see Table 1). The set of sequences observable by player i also includes personal histories of the form (h, a_i) : $H_i = H \cup \{(h, a_i) : h \in H, a_i \in A_i(h)\}$.

A CPS for player i over paths and beliefs of others is an array of probability measures $\beta_i = (\beta_i(\cdot|h_i))_{h_i \in H_i} \in [\Delta(Z \times \Delta_{-i})]^{H_i}$ satisfying conditions (2)-(3), which apply to the multistage setting as well. With this, also the notation on first- and second-order beliefs is the same as before: $\alpha_i \in \Delta_i^1 = \Delta^{H_i}(Z)$, $\beta_i \in \Delta_i^2 = \Delta^{H_i}(Z \times \Delta_{-i}^1)$. As before, α_i is first-order CPS derived from β_i when they appear in the same formula. The definitions of the equilibrium concepts SE and PSE can be applied to the multistage setting without modifications.

We distinguish between two extreme scenarios according to the behaviorally relevant periodization: In the **slow-play** scenario, stages correspond to periods, and the reference belief of player i at the beginning of period (stage) $t + 1$ is given by his belief at the beginning of period t . In the **fast-play** scenario, the different stages of the game occur in the same period and the relevant reference belief of player i in each stage t is given by his initial belief, that is, his belief at the root \emptyset .²⁴ Regardless of the scenario, we maintain the assumption that blame is continuous in beliefs and capped by frustration, $B_{ij}(h; \beta_i) \leq F_i(h; \alpha_i)$, and that $B_{ij}(h; \beta_i) = F_i(h; \alpha_i)$ in the case of simple anger.

6.1 Slow play

We start with this scenario because it allows for a relatively simple extension of the two-stage setting, with initial beliefs replaced by one-period-lagged beliefs: For any non-terminal history of the form $h = (\bar{h}, a)$ the frustration

²⁴Of course, in applications we may have intermediate cases, as in alternating-offer bargaining models where a period is composed of two stages. But the two extreme scenarios are sufficient to convey the main ideas.

of i conditional on h given α_i is

$$F_i(h; \alpha_i) = \left[\mathbb{E}[\pi_i | \bar{h}; \alpha_i] - \max_{a_i \in A_i(h)} \mathbb{E}[\pi_i | (h, a_i); \alpha_i] \right]^+.$$

(When $\bar{h} = \emptyset$ and $h = a^1$, we are back to the two-period formula.) With this, the decision-utility of action $a_i \in A_i(h^t)$ has the general form (4), where the blame functions $B_{ij}(h; \beta_i)$ are of the SA, ABB, or ABI type. Specifically: $B_{ij}(h; \beta_i) = F_i(h; \alpha_i)$ for SA, whereas the could-have-been blame, blaming deviations, and blaming intentions can be defined with straightforward adaptations of (8), (9), and (10) respectively; therefore we omit the details.

This extension of the two-stage setting has the stark feature that past frustrations do not affect current behavior. A more nuanced version of the slow-play scenario features a decaying effect of past frustrations, as in the following version of the SA decision-utility: Let $h^{\leq k}$ denote the prefix of length k of a history h of length $t \geq k$ (with $h^{\leq 0} = \emptyset$), then i 's **cumulated frustration at the beginning** of period $t + 1$ conditional on h given α_i is

$$\bar{F}_i(h; \alpha_i) = \sum_{k=1}^t (1+d_i)^{k-t} \left[\mathbb{E}[\pi_i | h^{\leq k-1}; \alpha_i] - \max_{a_i \in A_i(h^{\leq k-1})} \mathbb{E}[\pi_i | (h^{\leq k-1}, a_i); \alpha_i] \right]^+,$$

where $d_i \geq 0$ is a decay-rate parameter, and

$$u_i^{SA}(h, a_i; \alpha_i) = \mathbb{E}[\pi_i | (h, a_i); \alpha_i] - \theta_i \sum_{j \neq i} \bar{F}_i(h; \alpha_i) \mathbb{E}[\pi_j | (h, a_i); \alpha_i].$$

As $d_i \rightarrow \infty$, $\bar{F}_i(h; \alpha_i) \rightarrow F_i(h; \alpha_i)$ and we are back to the simple formulation where past frustrations do not affect current behavior.

A detail in modeling game forms becomes relevant in the slow play scenario: We have to explicitly allow for non-terminal histories after which no

player (not even chance) is active, such as history g in Figure D.




Figure D. Ultimatum Minigame with delayed reply.

Formally, at such histories there is only one feasible action profile, because each player has only one feasible action, to wait. In the two-periods setting this detail is irrelevant: If nobody is active at the root, play effectively starts (and ends) in the second period; if nobody is active at a^1 , a^1 can be modeled as a terminal history. But with more than two periods, having to wait may affect behavior.

Example 12 Consider Figure D. Suppose that Bob initially expects the fair offer f with positive probability. Then in period 2 after the greedy offer g he is frustrated, but cannot hurt Ann because he has to wait. In period 3, Bob’s lagged expectation has fully adapted downward, hence there is no “incremental frustration” in this period. According to our simple slow-play model, the frustration experienced by Bob in period 2 does not affect his decision utility in period 3: Bob fully “cools off” and behaves as-if selfish. Therefore the unique (polymorphic) SE outcome of the game is (g, w, y) , where w denotes waiting. In the model with decay the result is the same if d_b is large relative to θ_b : If Bob is initially certain of f , Bob’s cumulated frustration at (g, w) is $1/(1 + d_b)$ and he accepts if $\frac{1+d}{3\theta} > 1$. In this case, only (g, w, y) is a (polymorphic) SE; otherwise, there is a multiplicity of equilibrium outcomes including (f, w, y) . ▲

6.2 Fast play

When play is fast, all the stages belong to the same period, therefore the reference belief that determines player i 's frustration conditional on any history is i 's initially expected monetary payoff. Thus, i 's frustration at h given α_i is

$$F_i(h; \alpha_i) = \left[\mathbb{E}[\pi_i; \alpha_i] - \max_{a_i \in A_i(h)} \mathbb{E}[\pi_i | (h, a_i); \alpha_i] \right]^+.$$

This implies that *there cannot be any "cooling off"* due to reference-point acclimatization. Formally, histories where nobody (not even chance) is active play no role and can be deleted from the game form without affecting the analysis. For example, in the fast-play scenario, the ultimatum game form of Figure D is equivalent to the one of Figure B. The fast-play frustration formula can be plugged into the SA decision-utility function (6).

As for the ABB decision-utility, let us first extend the general property (7) of the blaming function B_{ij} :

$$B_{ij}(h; \alpha_i) = \begin{cases} 0, & \text{if } j \notin I(h') \text{ for all } h' \prec h, \\ F_i(a^1; \alpha_i), & \text{if } \{j\} = I(h') \text{ for all } h' \prec h. \end{cases} \quad (12)$$

In words, co-player j cannot be blamed if he was never active in the past, and he is fully blamed if instead he was the only active player. A relatively simple extension of could-have-been blame satisfies this property:

$$B_{ij}(h; \alpha_i) = \min \left\{ \left[\max_{h' \prec h, a'_j \in A_j(h')} \mathbb{E}[\pi_i | (h', a'_j); \alpha_i] - \mathbb{E}[\pi_i | h; \alpha_i] \right]^+, F_i(h; \alpha_i) \right\}. \quad (13)$$

We can follow a similar logic to extend ABI.

Remark 4 *If B_{ij} is defined by (13), then it satisfies (12).*

Proof Fix $h \in H$. First note that if j was never active before, then $A_j(h')$ is a singleton for each $h' \prec h$, hence the term in brackets of (13) is zero. Next suppose that i is frustrated at h and j was the only active player in the past. Then there must be some $\bar{h} \prec h$ such that j deviated from i 's expectations $\alpha_i(\cdot | \bar{h})$ for the first time, that is, \bar{h} is the shortest predecessor $h' \prec h$ such that $\alpha_j(a'_j | h') < 1$ for $(h', a'_j) \preceq h$. Such \bar{h} must have probability one according to the initial belief $\alpha_i(\cdot | \emptyset)$, thus

$\mathbb{E}[\pi_i|\bar{h}; \alpha_i] = \mathbb{E}[\pi_i; \alpha_i]$. Since $\max_{a'_j \in A_j(\bar{h})} \mathbb{E}[\pi_i|(\bar{h}, a'_j); \alpha_i] \geq \mathbb{E}[\pi_i|\bar{h}; \alpha_i]$, we have $\max_{a'_j \in A_j(\bar{h})} \mathbb{E}[\pi_i|(\bar{h}, a'_j); \alpha_i] \geq \mathbb{E}[\pi_i; \alpha_i]$. Therefore

$$\begin{aligned}
& \max_{h' \prec h, a'_j \in A_j(h')} \mathbb{E}[\pi_i|(h', a'_j); \alpha_i] - \mathbb{E}[\pi_i|h; \alpha_i] \\
& \geq \max_{a'_j \in A_j(\bar{h})} \mathbb{E}[\pi_i|(\bar{h}, a'_j); \alpha_i] - \mathbb{E}[\pi_i|h; \alpha_i] \\
& \geq \mathbb{E}[\pi_i; \alpha_i] - \mathbb{E}[\pi_i|h; \alpha_i] \\
& \geq \mathbb{E}[\pi_i; \alpha_i] - \max_{a_i \in A_i(h)} \mathbb{E}[\pi_i|(h, a_i); \alpha_i] = F_i(h; \alpha_i),
\end{aligned}$$

which implies $B_{ij}(h; \alpha_i) = F_i(h; \alpha_i)$ according to (13). ■

The following example illustrates our definition at work, elucidating a specific modeling choice:




Figure E. Multistage Ultimatum featuring Zoë.

Example 13 Consider the game form in Figure E (material payoffs are in alphabetical order). If Zoë chooses *In*, then Ann and Bob interact in an ultimatum minigame, but Zoë may instead exercise outside options and play (*Out*, *x*) or (*Out*, *y*). Zoë's payoffs equal Bob's, except following (*Out*, *y*) where a payoff transfer from Ann to Bob occurs, relative to (*Out*, *x*). Can

strategy profile $((In, x), f, n)$ be a SE under ABB? Given equilibrium beliefs, this is the case if $0 - \theta_b \cdot 1 \cdot 0 \geq 1 - \theta_b \cdot 1 \cdot 3$, or $\theta_b \geq 1/3$. The calculation involves Bob blaming Ann, not Bob blaming Zoë, because if Zoë switched from *In* to *Out* (thus implementing (Out, x) instead of *In*) this would not improve Bob's payoff. This reflects a non-obvious modeling choice. Our definition assesses blame on the basis of single-agent deviations from the realized path, but if Bob alternatively assessed blame on the basis of multi-agent deviations, including off-realized-path deviations, he would consider that Zoë could have played (Out, y) . She would then have increased Bob's payoff from 1 to 2, preventing his frustration of 1. If Bob's blame of Zoë were thus 1 then $((In, x), f, n)$ would be a SE under ABB if $0 - \theta_b \cdot 1 \cdot 0 \geq 1 - \theta_b \cdot 1 \cdot 3 - \theta_b \cdot 1 \cdot 1$, or $\theta_b \geq 1/4 \neq 1/3$. (This also shows that SE under ABB is not invariant with respect to coalescing sequential moves.) Finally, note that also $((In, y), f, n)$ is a SE under ABB in the fast-play scenario for $\theta_b \geq 1/4$, because at (In, g) Zoë would be blamed for not switching to *Out* (implementing (Out, y)); but it is a SE under ABB in the slow-play scenario for larger parameter values, $\theta_b \geq 1/3$, because Bob would be frustrated only in the third period, after (In, g) , and Zoë – who played in the first, could not be blamed. ▲

The single- vs. multi-agent deviation issue illustrated through this example can arise also in two-stage games (that involve simultaneous moves), but the point is clearer, and perhaps more relevant, in games with $n \geq 3$ stages. While noting the issue, we defend our chosen formulation thrice: It harmonizes well with the way we define rational play in this paper, where players optimize only locally (although in equilibrium they predict correctly and choose as planned). The (hinted at) alternative definition would be formally very convoluted. It is an open issue which formulation is empirically more relevant, so we stick with what is simpler.

6.3 Counterfactual anger and a unique SE in a hold-up problem

It is important to emphasize, in a separate section, that anger (and in fact emotions more generally) can shape behavior without occurring. If anger is anticipated, this may steer behavior at preceding histories down paths that render the sentiment counterfactual (see Remark 2). We already saw examples, *e.g.*, (f, n) is a SE in the Ultimatum mini-game, alongside (g, y) . Our next example, highlights how there may be circumstances where the

SE is *unique* and has that property. It furthermore illustrates a difference between fast and slow play.




Figure F. Hold-up.

Example 14 Modify the Ultimatum mini-game by adding an initial move for Bob, as seen in Figure F, to get an illustration of a hold-up problem (see Dufwenberg, Smith & Van Essen, 2013).²⁵ Under fast play, for each utility function seen so far,²⁶ if $\theta_b > 2/3$, there is a unique SE: Bob plays (r, n) and Ann plays f . To verify this, the key step is to check that the strategy profile where Bob plays (l, y) and Ann plays g is *not* an SE; if Bob initially expects \$1.5, off-path at (r, g) , he would be frustrated and hence would want to deviate to n . ▲

With slow play, by contrast, following r , with $\theta_b > 2/3$, SE play exhibits multiplicity, exactly as in the Ultimatum mini-game.

²⁵Bob and Ann face a joint business opportunity worth $(2, 2)$ via path (r, f) ; however, r involves partnership-specific investment by Bob, which Ann can exploit by choosing g to renege, etc. As always, we list payoffs by alphabetical order of players: (π_a, π_b) .

²⁶Except the blaming-unexpected-deviations version of ABB, which we did not define explicitly for fast play.

7 Discussion

Incorporating the effects of emotions in economic analysis is a balancing act. One wants to focus on sentiments that make empirical sense, but human psychology is multi-faceted and there is no unambiguous yardstick. Our chosen formulation provides a starting point for exploring how anger shapes economic interaction, and experimental or other evidence will help to assess empirical relevance and suggest revised formulas. We conclude by discussing topics that one way or another may be helpful for gaining perspective on, building on, or further developing our work. It is a mixture of commentary on chosen concepts, comparisons with related notions in the literature, and remarks about empirical tests.

Frustration The psychological evidence, cited in the Introduction, says that a player becomes frustrated when his goals are unexpectedly thwarted. We addressed but one aspect, concerning own material rewards. Cases 1-3 of the Introduction indicate the broad applied potential. Nevertheless, our focus is restrictive, in that one may imagine other sources of frustration. To see this, consider two more cases:

Case 4: In 2007 Apple launched its iPhone at \$499. Two months later they introduced a new version at \$399, re-priced the old model at \$299, and caused outrage among early adopters. Apple paid back the difference. Did this help long run profit?

Case 5: The 2008 TARP bank bail-out infuriated some US voters. Did this ignite the Tea Party/Occupy-WS movements?

These cases make sense, but our analysis does not address them directly (which is why we used other cases in the Introduction). In case 4, an early adopter is frustrated because he regrets he already bought, not because new information implies that his expected rewards drop. In case 5, an activist may be materially unaffected personally, yet frustrated because of unexpected perceived unfairness. Some meaningful ways to get frustrated are left for future research.

As regards the effects of frustration, we considered changes to a player's utility function but we neglected other plausible adjustments. Gneezy & Imas (2014) report data from an intriguing experiment involving two-player zero-sum material payoff games. In one game players gain if they are powerful, in

the other they are rewarded for being smart. Gneezy & Imas (2014) explore an added game-feature: before play starts, one subject may frustrate his opponent and force him to hang around in the lab to do boring tasks after the play ends. A thus frustrated player’s performance is enhanced when strength is beneficial (possibly from increased adrenaline flow), but reduced in the game where cool logic is called for (as if an angered player becomes cognitively impaired). Our model can capture the first consideration, but not the second. Specifically, we can let the consequences of play depend on beliefs as well as actions, *e.g.*, because emotions affect strength or speed (cf. Rauh & Seccia 2006); this ultimately translates into belief-dependent utility of actions that can capture the first effect. But, to capture the second effect, we would need a theory of endogenous cognitive abilities of boundedly rational players.

Valence and action-tendency Psychologists classify emotions in multiple ways. Two prominent aspects are **valence**, or the value the decision-maker associates with a sentiment, and **action-tendency**, or how behavior is shaped as the sentiment occurs. Both notions may, in principle, have bearing on anger. For example, anger may have negative valence, say if a frustration-laden life is taxing or decreases longevity. Perhaps such considerations steer people to avoid frustrations, say by not investing in the stock market. That said, the distinguishing feature of anger that psychologists stress concerns its action-tendency of aggression, not its valence. In developing our theory, we have exaggerated this, abstracting away from valence altogether while emphasizing aggression. This is reflected in the decision utility functions, which are shaped by current frustration, but there is no notion of valence. Anticipated future frustrations do not affect decision utility.²⁷

Blame We explored various ways a player may blame the frustration he experiences on others, yet more notions are conceivable. For example, with anger from blaming behavior player i ’s blame of j depends on what i believes he would truly get at counterfactual histories, not on the most he could get

²⁷In previous work we modeled another emotion, namely guilt – see, *e.g.*, Battigalli & Dufwenberg (2007), Chang, Smith, Dufwenberg & Sanfey (2011). To gain further perspective one may note that in that work our approach to valence and action-tendency was the opposite. Guilt may have valence (negative!) as well as action-tendency (say to engage in “repair behavior”; see *e.g.* Silver (2007). In modeling guilt we highlighted valence while neglecting action-tendency.

there were he not to vent any anger there. We would defend this modeling choice as a reflection of local agency and abstracting away from valence; i 's current agent views other agents as uncontrollable, and he has no direct care for their frustrations.

Another example relates to how we model anger from blaming intentions: i 's blame of j depends on β_i , his second-order beliefs. Recall that the interpretation concerns beliefs about beliefs about material payoffs. It does not concern beliefs about beliefs about frustration, which would be third-rather than second-order beliefs. Battigalli & Dufwenberg (2007), in a context which concerned guilt rather than anger, worked with such a third-order belief based notion of blame.

Our blame notions one way or another assess the *marginal* impact of other players. For example, consider a game where i exits a building while each $j \in I \setminus \{i\}$, unexpectedly to i , simultaneously hurls a bucket of water at i , who thus gets soaked. According to our blame notions, i cannot blame j as long as there are at least two hurlers. One could imagine alternatives where i blames, say, all the hurlers on the grounds that *collectively* they could have thwarted i 's misery.

Several recent experiments have explored interesting aspects of blame (Bartling & Fischbacher 2012, Gurdal *et al.* 2014, Celen, Schotter & Blanco 2014). Is a field of the economics-of-blame emerging? We wish to emphasize that our focus on blame is restricted to its relation to frustration only. Our paper is not about blame more generally, as of course there are many reasons besides frustration that may lead people to blame each other.²⁸

Kőszegi & Rabin Card & Dahl (2011) show that reports of domestic abuse go up when football home teams favored to win lose. They argue that this is in line with Kőszegi & Rabin's (2006, 2007) theory of expectations-dependent reference points. One way to think of our paper is as providing a formalization of that idea. Kőszegi & Rabin model the loss felt when a player gets less than he expected, which one may think of as a form disappointment with negative valence (cf. Bell 1985, Loomes & Sugden 1986). That account

²⁸For example, Celen *et al.* (2014) present a model where i asks himself how he would have behaved had he been in j 's position and had j 's beliefs, and where i blames j if j appears to be less generous to i than i would have been. This involves that one player may blame another even if he is not surprised/frustrated. Relatedly, one may imagine a model where players blame those they consider unkind, as defined in reciprocity theory (cf. the subsection below), again something independent of frustration.

per se does not imply that aggression follows, but it may seem natural to add such an angle. In other words, one may imagine a general model where disappointment/frustration has negative valence as well as an action-tendency of aggression. Kőszegi & Rabin’s model would be the modification that looks at the negative valence part only. Our model of simple anger focuses on the action tendency only, which is enough to capture the effect that Card & Dahl (2011) found.²⁹

Anger Management People aware of their inclination to be angry may attempt to ‘manage’ or ‘contain’ it. Our players anticipate how frustrations shape the behavior of themselves and others, and they may avoid or seek certain subgames because of that. However, there are other interesting related phenomena that we do not address: Can player i somehow adjust θ_i say by taking an “anger management class?” If so, would rational individuals want to raise, or to lower, their θ_i ? How might that depend on the game forms they play? These are potentially relevant questions related to how we have modeled action-tendency in this paper. Further issues would arise if we were to add some valence aspects of anger.

Rotemberg’s approach In a series of intriguing papers Rotemberg explores how consumer anger shapes firms’ pricing (2005, 2011), as well as interaction in ultimatum games (2008). He proposes (versions of) a theory in which players are slightly altruistic, and consumers/responders also care about their co-players’ degrees of altruism. Namely, they abruptly become very angry and punish a co-player whom they come to believe has an altruism parameter lower than some (already low) threshold. “One can thus think of individual i as acting as a classical statistician who has a null hypothesis that people’s altruism parameter is at least as large as some cutoff value. If a person acts so that i is able to reject this hypothesis, individual i gains ill-will towards this person” (Rotemberg 2008, p. 464).

On the one hand, as a literal statement of what makes people upset, this assumption does not match well our reading of the relevant psychology. Recall that frustration is anchored in goal-blockage, where individual are unexpectedly denied things they care about. Matters like “own payoff,”

²⁹Modeling details distinguish how we measure frustration from how Kőszegi & Rabin (2005, 2006) measure loss (*e.g.* concerning how we cap frustration at the highest attainable as opposed to actual payoff).

which is our focus, and “fairness” or “quality of past decisions,” which we have mentioned, come to mind; a co-player’s altruism being λ rather than $\lambda - \varepsilon$, where both λ and ε are tiny numbers, hardly does. On the other hand, it is impressive how well Rotemberg’s model captures the action in his data sets. It is natural to wonder whether our models could achieve that too. As regards behavior in ultimatum (and some other) games, there is already some existing evidence that is consistent with our modeling efforts; see the discussion in the final subsection below. Regarding pricing, we leave for empirical economists the task of exploring how our models might fare if applied to Rotemberg’s data sets.

Negative reciprocity Negative reciprocity (cf. Rabin 1993, Dufwenberg & Kirchsteiger 2004, Falk & Fischbacher 2006, Sebald 2010) joins anger as a form of motivation that can trigger hostile action. In some cases implications may be similar. However, anger and negative reciprocity differ in key ways and it is instructive to point out how. The following sketched comparison is with Dufwenberg & Kirchsteiger’s notion of sequential reciprocity equilibrium (SRE; refer to their article for formal definitions):

First, in the Hammering-One’s-Thumb game of Figure C, Andy may take it out on Bob if he is motivated by simple anger. If he were motivated by reciprocity, this could never happen: Andy’s kindness, since he is a dummy-player, equals 0, implying that a reciprocal Bob chooses as-if selfish. In this example reciprocity captures intuitions similar to the ABI concept, as perceived kindness assesses intentions similarly to how blame is apportioned.

Second, that analogy only carries so far, however. A player may be perceived as unkind even if he fails to hurt another, whereas under all our anger notions frustration is a prerequisite for hostile response. The following game

form illustrates:

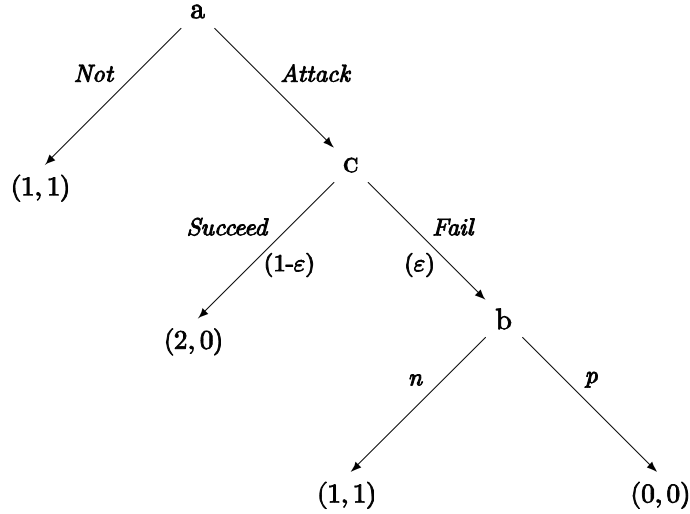


Figure G. Failed attack.

If b is asked to play then a 's attack failed. Under reciprocity theory (suitably augmented to allow incorporating a chance move; cf. Sebald 2010), b would deem a unkind, and – if sufficiently motivated by reciprocity – choose p in response. By contrast, under our anger concepts (SA, ABB, and ABI) b would not be frustrated, and since frustration is a prerequisite for hostility b would choose n .

Third, reciprocity allows for so-called “miserable equilibria,” where a player reciprocates expected unkindness before it occurs. For example, in the mini-ultimatum game of Figure B, (g, n) may be a SRE. Ann makes the greedy offer g despite believing that Bob will reject, because given her beliefs about Bob's beliefs, she perceives Bob as seeing this coming, which makes him unkind, so she punishes him by choosing g . This pattern of self-fulfilling prophecies of destructive behavior has no counterpart under either of our anger notions. Since Ann moves at the root, she cannot be frustrated, and hence, regardless of how prone to anger she may be in terms of anger sensitivity, she chooses as-if selfish.³⁰

³⁰Another example would be the hold-up game of Figure F. We gave conditions under which $((r, n), f)$ is the unique SE. By contrast, if Ann and Bob were motivated by reciprocity, $((l, n), g)$ and $((r, n), g)$ can be SRE, with miserable interaction, respectively, off and on the equilibrium path.

Fourth, with reference to our discussion of cooling-off effects in section 6, these have no counterpart in Dufwenberg & Kirchsteiger’s theory, which rather makes the same prediction in the games of Figures B and D. Reciprocal players do not cool off, they say things like “la vengeance est un plat qui se mange froid.”

Experimental testing Our models tell various stories of how interaction may play out when players prone to anger interact. It is natural to wonder about empirical relevance, and here experiments may be helpful. We would like to offer several related remarks.

First, several existing experimental studies provide evidence in favor of the notion that emotions drive behavior, and that many of them, and anger in particular, are generated from comparisons of outcomes with expectations: A number of studies find evidence for anger as the driving force behind costly punishment. A few papers rely on emotion-self reports: Pillutla & Murningham (1996) find that reported anger predicted rejections better than perceived unfairness in ultimatum games. Fehr & Gächter (2002) elicited self-reports of the level of anger towards free riders in a public goods game, concluding that negative emotions including anger are the proximate cause of costly punishment in the game. Other studies directly connect unfulfilled expectations and costly punishment in ultimatum games. Schotter & Sopher (2007) measure second-mover expectations in ultimatum games, concluding that unfulfilled expectations drive rejections of low offers. Similarly, Sanfey (2009) finds that psychology students who are told that a typical offer in the ultimatum game is \$4-\$5 reject low offers more frequently than students who are told that a typical offer is \$1-\$2.

A series of papers by Frans van Winden (with several coauthors) records both emotions and expectations in the power-to-take game (which resembles ultimatum games, but allows for partial rejections).³¹ These papers show that second-mover expectations about first-mover take rates in power-to-take games are a key factor in the decision to destroy income. Furthermore, they find that anger-like emotions are triggered by the difference between expected and actual take rates. In these experiments the difference between the take rate and the reported fair take rate is not significant in determining anger-like emotions, suggesting that deviations from expectations, rather than from

³¹Bosman & van Winden (2002), Bosman, Sutter & van Winden (2005), Reuben & van Winden (2008).

fairness benchmarks, drive both anger and the destruction of endowments in the games. Reuben and van Winden also argue that existing models of belief-dependent reciprocity miss an important trigger of anger by focusing on equilibrium predictions where actions are correctly anticipated. In our models anger arises only off the equilibrium path, addressing Reuben and van Winden’s point that anger is triggered by unfulfilled expectations.

A propos the cooling off effects discussed in section 6, Grimm & Mengel (2011) ran ultimatum game experiments with a treatment that forced responders to wait ten minutes before making their choice. Without delay less than 20% of low offers were accepted while 60–80% were accepted if the acceptance decision were delayed.

Finally, a literature in neuroscience connects expectations and social norms to study the neural underpinnings of emotional behavior. In Xiang, Lohrenz & Montague (2013), subjects respond to a sequence of ultimatum game offers whilst undergoing fMRI imaging. Unbeknownst to subjects, the experimenter controls the distribution of offers in order to manipulate beliefs. The authors find that rejections occur more often when subjects expect high offers relative to when they expect low offers. They make an important connection between norm violations and reward prediction errors from reinforcement learning, which are known to be the computations instantiated by the dopaminergic reward system. Xiang *et al.* note that “when the expectation (norm) is violated, these error signals serve as control signals to guide choices. They may also serve as the progenitor of subjective feelings.”

Going forward, it would be useful to develop tests specifically designed to target key features of our theory. For example, which version – SA, ABB, ABI – seems more empirically relevant, and how does the answer depends on context details (*e.g.*, is SA perhaps more relevant for tired subjects?). Some insights may again be gleaned from existing studies. For example, Gurdal *et al.* (2014) study games where an agent invests on behalf of a principal, choosing between a safe outside option and a risky alternative. If the latter is chosen, then it turns out that many principals punish the agent if and only if by chance a poor outcome is realized. This seems to indicate some empirical relevance of our ABB solution (relative to ABI). That said, Gurdal *et al.*’s intriguing design is not tailored to specifically test our theory (and beliefs and frustrations are thus not measured), so more work seems needed to draw clearer conclusions.

Finally, one must keep in mind that our models are abstractions. We theorize about the consequences of anger while neglecting myriad other obviously

important aspects of human motivation (say altruism, warm glow, inequity aversion, reciprocity, social status, or other emotions like guilt, disappointment, regret, or anxiety). Our models are not intended to explain every data pattern but rather to highlight the would-be consequences of anger, if anger were the only form of motivation at play (in addition to material payoff, of course). This statement may seem trivially obvious, but it has subtle implications for how to evaluate experimental work. To illustrate, consider again the Failed Attack game form in Figure G and suppose that in an experiment many subjects in player b 's position chose to punish (p). It would be silly to say that this constitutes a rejection of our theory (which predicts n rather than p), as what may obviously be going on is that one of the important forms of motivation that our theory deliberately abstracts away from is affecting subjects choices (presumably negative reciprocity, in line with our observations in the previous subsection). It would be sensible to ask, however, if those choices of p were in fact driven by anger (as might be measured by, *e.g.*, emotion self-reports, physiological activity, or both, as in Chang *et al.* 2011) and if they were (as opposed to being driven only by alternative motivations that we abstract away from, like negative reciprocity), then that would indicate that our theory could benefit from revision.

References

- [1] AKERLOF, R.J. (2014): "Anger and Enforcement," typescript, University of Warwick.
- [2] ANDERSON, E., AND D. SIMESTER (2010): "Price Stickiness and Customer Antagonism," *Quarterly Journal of Economics*, 125, 729-765.
- [3] ARNOLD, M.B. (1960): *Emotions and Personality*. New York: Columbia University Press.
- [4] ATTANASI, G., P. BATTIGALLI, AND E. MANZONI (2015): "Incomplete Information Models of Guilt Aversion in the Trust Game," *Management Science*, forthcoming.
- [5] AVERILL, J.R. (1983): "Studies on Anger and Aggression: Implications for Theories of Emotion," *American Psychologist*, 38, 1145-1160.

- [6] BARTLING, B., AND U. FISCHBACHER (2012): “Shifting the Blame: On Delegation and Responsibility,” *Review of Economic Studies*, 79, 67-87.
- [7] BATTIGALLI, P., AND M. DUFWENBERG (2007): “Guilt in Games,” *American Economic Review, Papers and Proceedings*, 97, 170-176.
- [8] BATTIGALLI, P., AND M. DUFWENBERG (2009): “Dynamic Psychological Games,” *Journal of Economic Theory*, 144, 1-35.
- [9] BATTIGALLI, P., AND M. SINISCALCHI (1999): “Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games,” *Journal of Economic Theory*, 88, 188-230.
- [10] BATTIGALLI, P., G. CHARNNESS, AND M. DUFWENBERG (2013): “Deception: The Role of Guilt,” *Journal of Economic Behavior & Organization*, 93, 227-232.
- [11] BATTIGALLI, P., A. DI TILLIO AND D. SAMET (2013): “Strategies and Interactive Beliefs in Dynamic Games,” in D. Acemoglu, M. Arellano, E. Dekel (eds.) *Advances in Economics and Econometrics: Theory and Applications*, Tenth World Congress, Vol. 1 Economic Theory (Econometric Society Monographs No. 49), Cambridge: CUP, pp 391-422.
- [12] BAUMEISTER, R. F. AND B. F. BUSHMAN (2007): “Angry Emotions and Aggressive Behaviors,” in G. Steffgen and M. Gollwitzer (Eds.) *Emotions and Aggressive Behavior*, Cambridge, MA: Hogrefe, pp. 61-75.
- [13] BELL, D. (1985): “Disappointment in Decision Making under Uncertainty,” *Operations Research*, 33, 1-27.
- [14] BERKOWITZ, L. (1989): “Frustration-Aggression Hypothesis: Examination and Reformulation,” *Psychological Bulletin*, 106, 59-73.
- [15] BOSMAN, R., M. SUTTER, AND F. VAN WINDEN (2005): “The Impact of Real Effort and Emotions in the Power-to-take Game,” *Journal of Economic Psychology*, 26, 407-429.
- [16] BOSMAN, R., AND F. VAN WINDEN (2002): “Emotional Hazard in a Power-to-take Experiment,” *Economic Journal*, 112, 147-169.

- [17] CARD, D., AND G. DAHL (2011): “Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior,” *Quarterly Journal of Economics*, 126, 103-143.
- [18] CARPENTER, J. AND P. MATTHEWS (2012): “Norm Enforcement: Anger, Indignation or Reciprocity,” *Journal of the European Economic Association*, 10, 555-572.
- [19] CELEN, B., A. SCHOTTER AND M. BLANCO (2014): “On Blame and Reciprocity: An Experimental Study,” typescript.
- [20] CHANG, L., A. SMITH, M. DUFWENBERG, AND A. SANFEY (2011): “Triangulating the Neural, Psychological, and Economic Bases of Guilt Aversion,” *Neuron*, 70, 560-572.
- [21] DOLLARD, J., L. DOOB, N. MILLER, O. MOWRER, AND R. SEARS (1939): *Frustration and Aggression*. Yale University Press, New Haven, CT.
- [22] DUFWENBERG, M., AND G. KIRCHSTEIGER (2004): “A Theory of Sequential Reciprocity,” *Games and Economic Behavior*, 47, 268-298.
- [23] DUFWENBERG, M., A. SMITH, AND M. VAN ESSEN (2013): “Hold-up: With a Vengeance,” *Economic Inquiry*, 51, 896-908.
- [24] FALK, A., AND U. FISCHBACHER (2006): “A Theory of Reciprocity,” *Games and Economic Behavior*, 54, 293-315.
- [25] FEHR, E., AND S. GACHTER (2002): “Altruistic Punishment in Humans,” *Nature*, 415, 137-140.
- [26] FRIJDA, N. H. (1986): *The Emotions*. Cambridge: CUP.
- [27] FRIJDA, N. H. (1993): “The Place of Appraisal in Emotion,” *Cognition and Emotion*, 7, 357-387.
- [28] GEANAKOPOLOS J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, 1, 60-79.

- [29] GNEEZY, U., AND A. IMAS (2014): “Materazzi Effect and the Strategic Use of Anger in Competitive Interactions,” *Proceedings of the National Academy of Sciences*, 111, 1334-1337.
- [30] GRIMM, V. AND F. MENGEL (2011), “Let Me Sleep on It: Delay Reduces Rejection Rates in Ultimatum Games,” *Economics Letters*, 111, 113-115.
- [31] GURDAL, M., J. MILLER AND A. RUSTICHINI (2014): “Why Blame?” *Journal of Political Economy*, 121, 1205-1246.
- [32] KŐSZEGI, B., AND M. RABIN (2006): “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121, 1133-1166.
- [33] KŐSZEGI, B., AND M. RABIN (2007): “Reference-Dependent Risk Attitudes,” *American Economic Review*, 97, 1047-1073.
- [34] KREPS D., AND R. WILSON (1982): “Sequential Equilibrium,” *Econometrica*, 50, 863-894.
- [35] LOOMES, G., AND R. SUGDEN (1986): “Disappointment and Dynamic Consistency in Choice under Uncertainty,” *Review of Economic Studies*, 53, 271-282.
- [36] PASSARELLI, F., AND G. TABELLINI (2013): “Emotions and Political Unrest,” Bocconi University, IGIER Working Paper 474.
- [37] PILLUTLA, M., AND K. MURNINGHAM (1996): “Unfairness, Anger, and Spite: Emotional Rejections of Ultimatum Offers,” *Organizational Behavior and Human Decision Processes*, 68, 208-224.
- [38] POTEHAL, M., C. SPIELBERGER AND G. STEMMLER (eds.) (2010): *International Handbook of Anger: Constituent and Concomitant Biological, Psychological, and Social Processes*. Springer Verlag.
- [39] RABIN, M. (1993): “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, 83, 1281-1302.
- [40] RAUH, M. AND G. SECCIA (2006): “Anxiety and Performance: A Learning-by-Doing Model,” *International Economic Review*, 47, 583-609.

- [41] REUBEN, E., AND F. VAN WINDEN (2008): “Social Ties and Coordination on Negative Reciprocity: The Role of Affect,” *Journal of Public Economics*, 92, 34-53.
- [42] ROTEMBERG, J. (2005) “Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy,” *Journal of Monetary Economics*, 52, 829-852.
- [43] ROTEMBERG, J. (2008): “Minimally Acceptable Altruism and the Ultimatum game,” *Journal of Economic Behavior and Organization*, 66, 457-476.
- [44] ROTEMBERG, J. (2011): “Fair Pricing,” *Journal of the European Economic Association*, 9, 952-981.
- [45] SANFEY, A. (2009): “Expectations and Social Decision-making: Biasing Effects of Prior Knowledge on Ultimatum Responses,” *Mind and Society* 8, 93-107.
- [46] SCHOTTER, A., AND B. SOPHER (2007): “Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach,” *Games and Economic Behavior*, 58, 365-393.
- [47] SEBALD, A. (2010): “Procedural Concerns and Reciprocity,” *Games and Economic Behavior*, 68, 339-352.
- [48] SELTEN, R. (1975): “Re-Examination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25-55.
- [49] SILFVER, M. (2007): “Coping with Guilt and Shame: A Narrative Approach,” *Journal of Moral Education*, 36, 169-183.
- [50] SMITH, A. (2009): “Belief-Dependent Anger in Games,” typescript, University of Arizona.
- [51] XIANG, T., T. LOHRENTZ AND R. MONTAGUE (2013): “Computational Substrates of Norms and Their Violations during Social Exchange,” *Journal of Neuroscience*, 33, 1099-1108.