

Dulleck, Uwe; Löffler, Andreas

Article

“; games

Games

Provided in Cooperation with:

MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Dulleck, Uwe; Löffler, Andreas (2021) : “; games, Games, ISSN 2073-4336, MDPI, Basel, Vol. 12, Iss. 1, pp. 1-12, <https://doi.org/10.3390/g12010005>

This Version is available at:

<https://hdl.handle.net/10419/257488>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.



You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

μ - σ Games

Uwe Dulleck^{1,2,†}  and Andreas Löffler^{3,*,†} 

¹ Centre for Behavioural Economics, Society and Technology (BEST), Queensland University of Technology, Brisbane, QLD 4000, Australia; uwe.dulleck@qut.edu

² Crawford School of Public Policy, Australian National University, Canberra, ACT 4000, Australia

³ Freie Universität Berlin, 14195 Berlin, Germany

* Correspondence: al@wacc.de

† These authors contributed equally to this work. The authors would like to thank Adam Clements for comments and feedback.

Abstract: Risk aversion in game theory is usually modeled using expected utility, which was criticized early on, leading to an extensive literature on generalized expected utility. In this paper we are the first to apply μ - σ theory to the analysis of (static) games. μ - σ theory is widely accepted in the finance literature; using it allows us to study the effect on uncertainty endogenous to the game, i.e., mixed equilibria. In particular, we look at the case of linear μ - σ utility functions and determine the best response strategy. In the case of 2×2 and $N \times M$ games, we are able to characterize all mixed equilibria.

Keywords: μ - σ utility; game theory; mixed strategies; equilibrium



Citation: Dulleck, U.; Löffler, A. μ - σ Games. *Games* **2021**, *12*, 5.
<https://doi.org/10.3390/g12010005>

Received: 4 August 2020

Accepted: 8 January 2021

Published: 12 January 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

It is a well established fact in economic research that people differ in their attitudes towards risk and uncertainty. In game theoretic models of strategic interaction, the main tool to capture risk and uncertainty is expected utility theory [1]. This model was criticized by [2,3], and that led to an extended literature on generalized expected utility models that are able to accommodate the identified paradoxes. To name just few, [4,5] proposed models of rank-dependent utility relaxing the independence axiom; [6] relaxed the reduction of compound lotteries axiom to explain the Ellsberg paradox. Contributions to this literature usually start out by relaxing axioms that are the basis for von Neumann and Morgenstern's model, and show that by relaxing these axioms one can find a generalized model of expected utility while keeping most of the attractiveness of the standard model, and at the same time, enriching the framework to allow for seemingly irregular behavior with respect to the standard model.

One idea that comes to mind is the μ - σ theory. Its equilibrium equivalent, the CAPM—capital asset pricing model—describes the relationship between systematic risk and expected returns for assets traded on markets, and has wide applications in finance—as every textbook (see, among few, ([7] (Section 9.1)) and many papers (a still contemporary and in addition critical overview is [8]) show. CAPM, with its statements on optimal portfolio choice, is one of the most widely used models in finance today. Therefore, it seems worthwhile to understand its effects on game theory. We follow the finance approach and capture risk not using expected utility, but using μ - σ utility (see [9]).

Sometimes, μ - σ is seen as a special case of quadratic utility functions, of the form $u(x) = ax - bx^2$, with x being the payoffs of a lottery or an asset. We do not restrict ourselves to this special case but will deal with a general form $V(\mu, \sigma^2)$. These utility functions are considered today not as a special case of expected utility, but as an entire

different coverage of risk.¹ This approach will require using monetary (material) instead of utility based payoffs in the game.²

Most of the above-mentioned generalizations in game theory can be formalized using Choquet integrals.³ Notice that our approach cannot be described by Choquet integrals, but modern finance theory is based it on its own axiomatization.⁴

In game theory, the linear formulation of expected utility (with respect to probabilities) does not allow one to capture preferences over uncertainty endogenous to the game. A mixed strategy of a player, whether interpreted as the belief of another player or a real randomization, causes uncertainty for players. In the standard model, due to the linear formulation, this uncertainty is treated like it would be under the assumption of risk neutrality, i.e., in fact ignored. Looking at μ - σ utility, the circumstances are different—since probabilities enter the variance σ^2 (due to the fact that $\sigma^2 = E[x^2] - E[x]^2$ with x being payoffs). Hence, any mixed strategy of a player will now cause real uncertainty that will not be disregarded. Thus, in many cases where mixed-strategy equilibria exist in games assuming expected utility, they do not exist when one assumes μ - σ utility. That players tend to avoid mixed strategies in economic experiments can be seen, for example, in the Stag-Hunt game (see [19,20]).⁵

In this article we discuss how equilibrium predictions change. As we concentrate on two-player games, we interpret a mixed strategy as a real randomization by a player and not as a belief of the composition of population from which the other player is chosen randomly, with this player then choosing a pure strategy. One important, additional aspect of the interpretation of a mixed strategy as real randomization is that in the case of mixed equilibria, a player's own strategy now affects his or her utility, even though it does not change the expected value of the payoff. One may interpret this in the context of repeated play as a cost to changing ones actions in different rounds of the game. In this context, the literature on ambiguity aversion [22,23] comes to mind. Our approach is one of ex-ante utility—i.e., randomizing over ones actions is costly as it increases the variance. Taking a population interpretation of mixed strategy equilibria and applying payoff dominance as selection criterion [24] shares with our approach that for coordination games, mixed strategies are not selected as mixing reduces expected payoffs—whereas with the μ - σ approach there is a direct cost of randomization in form of the increased variance.

To refer back to finance theory where expected utility theory and μ - σ theory would argue that a portfolio does better than a single investment, a similar result in the application of μ - σ theory to game theory will not hold! In static games based on standard expected utility, reducing the variance of (monetary) payoffs is not important because the expected value of the utility payoffs stay the same; but μ - σ theory will indicate that a pure strategy, i.e., choosing one action instead of randomizing, is preferable. By explicitly capturing variance caused by an agent's own strategy choice, we may provide some reasoning as to why experimental players refrain frequently from using mixed strategies.

To illustrate our argument, consider the simple μ - σ utility function $V(\mu, \sigma^2) = \mu - \frac{r}{2}\sigma^2$. This utility function is linear in expectation and variance, and therefore is sometimes called "linear utility", although linearity here does not refer to the material payoffs or the probabilities of the players. In a typical μ - σ^2 diagram, any indifference curves are upward sloping, which follows from the fact that a higher variance needs to be compensated by a higher expected value; in our case, the indifference curves are straight lines with slope $\frac{\partial V}{\partial \mu} / \frac{\partial V}{\partial \sigma^2} = \frac{r}{2}$; for details, see ([12] (p. 426)). We focus on this particular utility function

¹ See, for example, [10–12].

² References [13–16] also consider monetary or material games.

³ See [17].

⁴ See [18] for an axiomatic foundation of μ - σ utility functions based on preference axioms. That μ - σ utility functions are not a special case of a Choquet integral using some capacity can easily be seen: Choquet integrals are homogenous of degree one, a feature that many μ - σ functions (for example, μ - σ^2) do not possess.

⁵ However, notice that [21] found very high rates of randomization in their experiments.

because it plays a major role in finance: Every investor with this utility function will exhibit a behavior known from “constant absolute risk aversion” in the expected utility framework (see ([12] (property 5))). This utility function serves as a starting point for understanding our idea.

We now consider a game where a player chooses a possibly mixed strategy given the possibly mixed strategy profile of other player(s); then the player faces a lottery where material payoffs depend on the specified game, and probabilities depend on the strategy profile. For a given strategy of the other player(s), each strategy of a player represents a point in the μ - σ diagram given the expected value and the implied variance of the strategy; i.e., for any strategy α this point is determined by $\mu = E(\alpha)$ and $\sigma^2 = \text{Var}(\alpha)$. We will show that mixing between two strategies of player i with the same utility (see Figure 1, where these strategies are denoted α_i and α'_i) actually leads to a of utility as the variance *increases*. This is in sharp contrast to the usual "egg shaped" efficient frontier seen in almost every textbook in finance, where mixing decreases the variance and therefore contributes to an increase in utility.

One main difference relative to other applications of non-generalized or generalized expected utility functions to game theory (see, for example, [25]) is that terminal node utilities of players are now dependent on how this terminal node is reached. In the case of random events, whether due to moves by nature or mixed strategies of one of the players, payoffs that are identical with respect to their material (or monetary) payoffs give rise to differences in utility under the μ - σ paradigm. That terminal node utility may depend on endogenous aspects of other players' behavior has recently received some detailed attention under the heading of psychological game theory (see, for example, [15,16,26]) where beliefs of players matter for the utility payoffs of players). μ - σ utility in games implies that strategies and the history of play, whenever random events are involved, affect the utility.

We proceed by first defining games based on μ - σ utility functions and study static 2×2 and $N \times N$ games based on linear utility functions. We then discuss non-linear μ - σ utility functions.

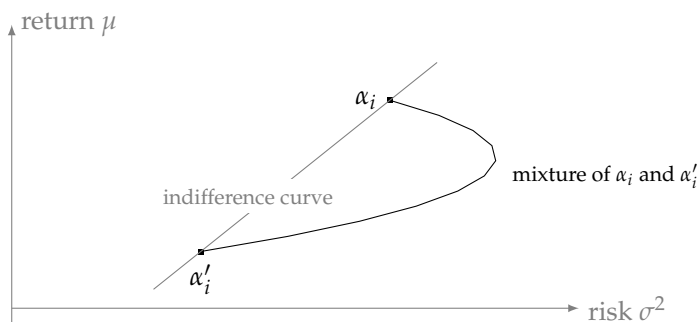


Figure 1. Mixing two lotteries α_i and α'_i with the same utility decreases the utility in a μ - σ game.

2. Definition of Static μ - σ Games

To understand the applicability of μ - σ utility functions to games, we start by analyzing static games with complete information. We consider games with a finite set N of players and nonempty and finite sets A_i ($i \in N$) of actions. Any profile of pure strategies $a \in \prod_{i=1}^N A_i$ will provide player i with a material (not utility) payoff $u_i(a)$; we use the notation $a = (a_1, \dots, a_N)$.

We consider mixed strategies—that is, elements of $\prod_i \Delta(A_i)$. Let α be a profile of mixed strategies, that is, a vector $\alpha = (\alpha_1, \dots, \alpha_N)$ with $\alpha_i \in \Delta(A_i)$. Furthermore, $\alpha_i(a_i)$ with $a_i \in A_i$ is the probability that player i will play the pure strategy $a_i \in A_i$ and $\alpha(a) = \prod_i \alpha_i(a_i)$ for $a \in \prod_i A_i$.

Player i can expect the following material payoff from strategy combination α

$$E[\alpha] := \sum_{a \in \times_i A_i} \alpha(a) u_i(a) \tag{1}$$

and a variance of

$$\text{Var}[\alpha] := \sum_{a \in \times_i A_i} \alpha(a) u_i^2(a) - E^2[\alpha] \tag{2}$$

For ease of notation we use $\mu = E[\alpha]$ and $\sigma^2 = \text{Var}[\alpha]$.

Definition 1. A μ - σ game is a game where the utility of player from strategy combination α is given by a μ - σ utility function

$$U_i(\alpha) = V_i(\mu, \sigma^2).$$

V_i is strictly increasing in the first variable and strictly decreasing in the second variable, and is strictly quasiconcave in μ and σ^2 .

We assume strict quasiconcavity to ensure uniqueness of the solution to classical maximization problems in finance. Relaxing this assumption most likely does not alter our results but makes the arguments very tedious.⁶

We first analyze the case of two players. Furthermore, we assume that the utility function is of the following simple linear form:

$$V(\mu, \sigma^2) = \mu - \frac{r}{2} \sigma^2$$

with r being the parameter for the strength of the variance aversion.⁷ We refer to a utility function of this form as linear utility.⁸

In the next paragraphs we analyze the effect of this utility model in well known examples of the literature on game theory. This allows us to show that in μ - σ games a Nash equilibrium does not always exist.

3. First Results for μ - σ Games

3.1. Best Response with Linear Utility

Compared to standard game theory, μ - σ games may have different sets of equilibria. In standard game theory mixed strategies, i.e., strategies that randomize over actions that lead to the same expected (material) utility, yield the same utility payoff for a player, due to the the linearity in probabilities assumed by the expected utility framework. For μ - σ games the randomization of a mixed strategy comes at a price.

Given the behavior of the other player(s) in the game, the following maximization problem determines the best response of a player with μ - σ utility.

$$\begin{aligned} & \max_{\alpha_i} \sum_{a \in \times_i A_i} \alpha(a) u_i(a) - \frac{r}{2} \left(\sum_{a \in \times_i A_i} \alpha(a) u_i^2(a) - E^2[\alpha] \right) \\ & = \max_{\alpha_i} \sum_{a \in \times_i A_i} \alpha(a) \left(u_i(a) - \frac{r}{2} u_i^2(a) \right) - \frac{r}{2} \left(\sum_{a \in \times_i A_i} \alpha(a) u_i(a) \right)^2 \end{aligned}$$

⁶ See the discussion in [10].

⁷ This is analogue to the definition of the Arrow-Pratt measure, see [10,12].

⁸ Note, this utility function shares with CARA expected utility functions the characteristic that the preference overall risk is independent of the wealth or income level.

This is a quadratic equation in α_i , where the coefficient on the quadratic term $a_i^2(a_i)$ is always negative.

Lemma 1 (best response with mixed strategy). *In equilibrium a player with μ - σ utility does not choose a mixed strategy, unless all actions chosen with positive probability are characterized by the same expected value and the same variance—i.e., they are characterized by the same point in the μ - σ^2 diagram.*

Proof. We show first that all convex combinations of two (mixed) strategies lie on a convex curve in a μ - σ diagram. Let α_i and α'_i be two strategies of player i . Expected value and variance given the strategies of other players can be denoted as

$$(E[\alpha_i], \text{Var}[\alpha_i]) \text{ and } (E[\alpha'_i], \text{Var}[\alpha'_i]).$$

Let a convex mixture choose $\lambda < 1$ times to play α_i and $1 - \lambda$ times to play α'_i ; we refer to this strategy as λ . The expected value can then be calculated following (1):

$$\begin{aligned} E[\lambda] &= \lambda E[\alpha_i] + (1 - \lambda)E[\alpha'_i] \\ &= E[\alpha'_i] + \lambda(E[\alpha_i] - E[\alpha'_i]). \end{aligned} \tag{3}$$

and the variance is given as

$$\begin{aligned} \text{Var}[\lambda] &= \sum_{a \in \times_i A_i} (\lambda \alpha_i + (1 - \lambda)\alpha'_i) \alpha_{-i}(a) u_i^2(a) - E^2[\lambda] \\ &= \lambda \sum_{a \in \times_i A_i} \alpha_i \alpha_{-i}(a) u_i^2(a) + (1 - \lambda) \sum_{a \in \times_i A_i} \alpha'_i \alpha_{-i}(a) u_i^2(a) \\ &\quad - (\lambda E[\alpha_i] + (1 - \lambda)E[\alpha'_i])^2 \\ &= \lambda (\text{Var}[\alpha_i] + E^2[\alpha_i]) + (1 - \lambda) (\text{Var}[\alpha'_i] + E^2[\alpha'_i]) \\ &\quad - (\lambda E[\alpha_i] + (1 - \lambda)E[\alpha'_i])^2 \\ &= \lambda \text{Var}[\alpha_i] + (1 - \lambda)\text{Var}[\alpha'_i] + \lambda(1 - \lambda)(E[\alpha_i] - E[\alpha'_i])^2. \end{aligned} \tag{4}$$

In the case that $E[\alpha_i] \neq E[\alpha'_i]$

$$\begin{aligned} \text{Var}[\lambda] &= \frac{E[\lambda] - E[\alpha'_i]}{E[\alpha_i] - E[\alpha'_i]} \text{Var}[\alpha_i] + \left(1 - \frac{E[\lambda] - E[\alpha'_i]}{E[\alpha_i] - E[\alpha'_i]}\right) \text{Var}[\alpha'_i] + \\ &\quad + \frac{E[\lambda] - E[\alpha'_i]}{E[\alpha_i] - E[\alpha'_i]} \left(1 - \frac{E[\lambda] - E[\alpha'_i]}{E[\alpha_i] - E[\alpha'_i]}\right) (E[\alpha_i] - E[\alpha'_i])^2 \end{aligned}$$

and the second derivative is given as

$$\frac{d^2 \text{Var}[\lambda]}{dE[\lambda]^2} = -2.$$

Therefore, any convex combinations of two strategies lie on a convex curve. Figure 2 illustrates such curves for three different combinations of strategies.

Our result follows from the observation that indifference curves of a player in the μ - σ diagram are given by straight lines between two strategies, given the linear utility function

$\mu - \frac{r}{2}\sigma^2$. For this reason any convex combination of strategies must be worse than either of the strategies that are combined. Even if a player is indifferent between two strategies α_i and α'_i , he will see any mixture between these strategies as inferior. \square

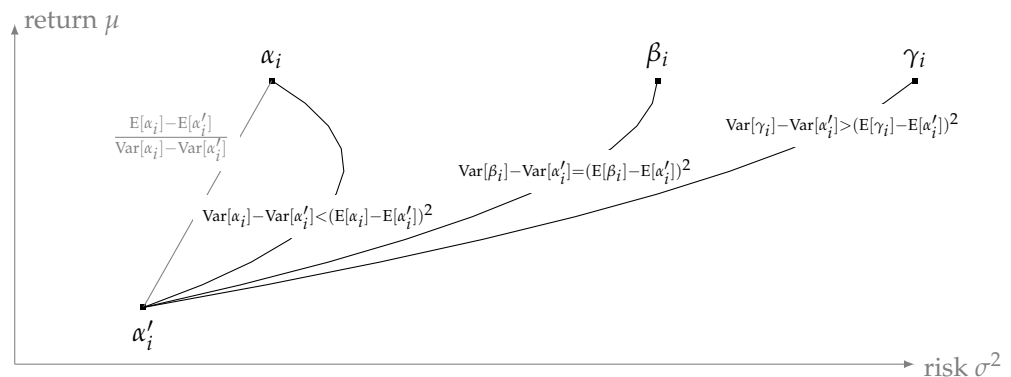


Figure 2. Mixtures of three pairs of strategies (α_i, α'_i) , (β_i, α'_i) and (γ_i, α'_i) in the μ - σ diagram.

This implies an equilibrium in mixed strategies only exists if both strategies are represented by the same point in the μ - σ diagram—i.e., they do not only have the same expected material payoff but also lead to the same variance.

Although this lemma seems to be related to Lemma 1 in [14], our model differs from theirs in one substantial point. The set of all pure actions in [14] is convex, which is not the case in static games with a finite number of pure strategies. Similarly, [13] looked at games where the players violate von Neuman and Morgenstern’s independence axiom, and assumed that the preferences (in terms of payments) were quasiconcave. Again, our paper differs from that work because (in terms of payments) μ - σ utility functions need not be quasiconcave.⁹

We next study the implication of this with respect to the best response towards a pure strategy and when an equilibrium in mixed strategies exist.

Lemma 2 (best response given pure strategies of another player). *The best response to a pure strategy is a pure strategy, unless the material payoff of the player under consideration is the same for a set of at least two actions. Any (best response) mixed strategy can only randomize over this set of actions.*

Proof. Referring again to Figure 2. Given that the other player chooses a pure strategy, all strategies of the player under consideration will be a point on the μ -axis. This implies that the point higher on the axis will be chosen unless two strategies lead to exactly the same material payoff. \square

3.2. 2×2 Games with Linear Utility

To answer the question when mixed strategy equilibria exist in μ - σ games, we start with the case of 2×2 games. The game we consider is given in Figure 3. We state our result for player 1 (the row player) and denote by q the probability that player 2 (the column player) chooses left. The following lemma characterizes the necessary condition for a best mixed strategy best response in comparison to the condition assuming standard expected utility theory.

⁹ Using suitable numbers, one can already show that $\mu - \sigma^2$ has upper contour sets that are not convex.

	left	right
up	a, ·	b, ·
down	c, ·	d, ·

Figure 3. Material payoffs for best response function.

Lemma 3 (best response in 2×2 games). *The best response in any 2×2 game is a mixed strategy if and only if, (a) the usual condition of expected utility game theory holds,*

$$0 = (a - c)q + (b - d)(1 - q), \tag{5}$$

and (b) the following condition is true:

$$a + c = b + d. \tag{6}$$

This lemma shows that 2×2 μ - σ games do not have more mixed equilibria than an equivalent standard 2×2 game. Being a mixed strategy equilibrium in the standard game is a necessary but not sufficient condition for being an equilibrium in the equivalent μ - σ game. The second condition (6) has to be fulfilled as well; therefore, many μ - σ games will not have any equilibrium.

Condition (6) has an insightful interpretation. It requires that, given any pure strategy of the other player, the sum over all material payoffs that the player is able to achieve over all his strategies is constant, i.e., independent of the pure strategy that is chosen given the randomization of the other player. Regardless the other player’s choice, it is only the slice of the cake and not the size of the cake that is determined by the player’s own actions.

Proof. We apply Lemma 1, which implies for the 2×2 game that expected value and variance have to be the same for any variation in the probability p of player 1 to choose up. This leads to the following two conditions:

$$\begin{aligned} E[p] &= apq + bp(1 - q) + c(1 - p)q + d(1 - p)(1 - q) = \text{const} \\ \text{Var}[p] &= a^2pq + b^2p(1 - q) + c^2(1 - p)q + d^2(1 - p)(1 - q) - (E[p])^2 = \text{const}. \end{aligned}$$

The constant expected value implies

$$(a - c)q + (b - d)(1 - q) = 0.$$

A non-degenerate mixed equilibrium requires $0 < q < 1$, and thus $a \neq c$ and $b \neq d$. Simplifying the condition under constant variance by solving for $E[p]$ and calculating the first-order condition p , gives us the second condition, so it follows as

$$(a^2 - c^2)q + (b^2 - d^2)(1 - q) = 0.$$

Combining both conditions gives us the second condition for the existence of a mixed equilibrium $a + c = b + d$. \square

We next characterize the games where mixed equilibria do survive.

Theorem 1 (Mixed Equilibria in 2×2 games). *A mixed equilibrium in 2×2 μ - σ games with linear utility functions exists if and only if*

- (i) *The candidate for equilibrium is a mixed equilibrium of a standard (expected utility) game with utility payoffs equal to the monetary payoffs of the μ - σ game;*
- (ii) *For each strategy of the other player, the sum of monetary payoffs of a player is the same for the strategies available to the player.*

Proof. The two conditions imply the existence of a mixed strategy equilibrium follows directly from Lemma 3, in particular Equations (5) and (6). \square

In the following we discuss special cases. We start by analyzing zero-sum games.

Theorem 2 (2×2 zero-sum games). *The only 2×2 - μ - σ zero-sum game with an equilibrium in non-degenerate mixed strategies is matching pennies.*

Proof. We denote the material payoffs as given by Figure 4.

Given the previous results, we know that the following two conditions have to be fulfilled in a mixed strategy equilibrium

$$a + c = b + d \tag{7}$$

$$\alpha + \beta = \gamma + \delta \tag{8}$$

Given that we study zero-sum games we know

$$a + \alpha = b + \beta = c + \gamma = d + \delta = C.$$

Substituting the last equation into the previous two, gives us

$$C - \alpha + C - \gamma = C - \beta + C - \delta.$$

Adding this equation to (8), yields $\beta = \gamma$ and $b = c$ and $a = d$ and $\alpha = \delta$; therefore, Figure 5 represents this game. \square

	<i>left</i>	<i>right</i>
<i>up</i>	a, α	b, β
<i>down</i>	c, γ	d, δ

Figure 4. Material payoffs in the game of Theorem 2.

	<i>left</i>	<i>right</i>
<i>up</i>	$a, C - a$	$b, C - b$
<i>down</i>	$b, C - b$	$a, C - a$

Figure 5. Matching pennies—the only μ - σ zero-sum game with (non-degenerate) mixed strategies.

Our next result concerns non-zero-sum games. It is well-known—despite disturbing results in μ - σ theory—that a portfolio with higher payoffs is not necessarily preferred by an investor (preferences need not be monotone).¹⁰ Therefore, it is not clear that dominated strategies in μ - σ games cannot be equilibrium strategies. In 2×2 games we can show the following result.

Theorem 3 (2×2 games with dominated strategies). *If a strategy in a 2×2 game is dominated in monetary payoffs, then no mixed equilibrium of the 2×2 μ - σ games exists.*

Proof. This follows immediately from the fact any equilibrium of the μ - σ game must be an equilibrium in the equivalent expected utility framework. In expected utility, games dominated strategies never receive a positive probability weight. \square

While this result seems to be obvious at first, it is less so if one considers that a player’s monetary payoff choosing the—in monetary payoffs—dominant action may lead to a higher variance than the dominated action over compensating for the loss in payoff. As the result shows, this cannot be the case. These results have a set of implications that are noteworthy:

¹⁰ This particular feature of μ - σ preferences is well known; see, for example, [27]. The fact that the CAPM as an application of μ - σ is still used today shows that the non-monotonicity is not regarded as a major problem in finance.

1. Only coordination games and games without a pure strategy equilibrium in the expected utility framework can have mixed strategy equilibria. This follows from the observation that solving (5) for q and substituting (6). Given $q \in (0, 1)$, either one action dominates the other or players prefer payoffs in two diagonal corners. Theorem 3 rules out the former, leaving the cases where both players either prefer the same two diagonal corners (coordination games) or they prefer different corners—games without a pure strategy equilibrium in standard games.
2. Battle of sexes μ - σ games do not have a mixed strategy equilibrium unless players—in case of miscoordination—receive an additional payoff equal to the difference in their payoffs between the preferred and the alternative equilibrium. This can be seen from Figure 6. A mixed strategy equilibrium exists iff (6) is satisfied. This is equivalent to $a = b + c$.

	left	right
up	a,b	c,c
down	0,0	b,a

Figure 6. Battle of sexes with mixed strategy equilibrium ($a > b$).

3.3. $N \times M$ Games with Linear Utility

Mixed equilibria in 2×2 μ - σ games only exist if an additional constraint on payoffs holds to ensure that these strategies are a best response. In this section we show that for $N \times M$ games a mixed equilibrium only exists if the game is degenerate. To show this, we show that any best response avoids mixing unless the material payoff for this player are constant over all possible outcomes of the game.

Theorem 4 (best response with $N \times M$ games). *In an $N \times M$ μ - σ game ($M, N > 2$) in any mixed equilibrium, players randomize at most over two pure strategies, unless the payoffs to the player are constant (independent of his choice).*

Proof. Again, we study the best responses of players. From Lemma 1 we know that any action that may be chosen by a player will have to be represented by the same point in the μ - σ diagram. To illustrate our argument, let us assume that a player randomizes over 3 actions, while the other player randomizes only over two. To show this, consider the 3×2 game with material payoffs given in Figure 7.

Let p_1, p_2 , and $1 - p_1 - p_2$ be the probabilities that the player chooses up, middle, and down respectively. The condition for a constant expected value in this case is

$$\text{const} = E[p_1, p_2] = ap_1q + bp_1(1 - q) + cp_2q + dp_2(1 - q) + e(1 - p_1 - p_2)q + f(1 - p_1 - p_2)(1 - q).$$

This is a condition on two variables, which gives us two constraints:

$$\begin{aligned} 0 &= (a - e)q + (b - f)(1 - q) \\ 0 &= (c - e)q + (d - f)(1 - q). \end{aligned}$$

If one combines both, they imply also $(a - c)q + (b - d)(1 - q) = 0$. Any non-degenerate mixed equilibrium implies that $q \neq 0$ which immediately implies $a \neq e \neq c$.

Furthermore the variance needs to be constant:

$$\text{const} = \text{Var}[p_1, p_2] = a^2p_1q + b^2p_1(1 - q) + c^2p_2q + d^2p_2(1 - q) + e^2(1 - p_1 - p_2)q + f^2(1 - p_1 - p_2)(1 - q) - (E[p_1, p_2])^2.$$

This is a condition on two variables, and using derivatives it can be reduced to two equations

$$0 = (a^2 - e^2)q + (b^2 - f^2)(1 - q)$$

$$0 = (c^2 - e^2)q + (d^2 - f^2)(1 - q)$$

and thus $0 = (a^2 - c^2)q + (b^2 - d^2)(1 - q)$. Solving implies

$$a + e = b + f$$

$$c + e = d + f$$

$$a + c = b + d$$

These three equations imply that $a = c = e$ and $b = d = f$, which contradicts the condition for the constant expected value. □

	<i>left</i>	<i>right</i>
<i>up</i>	a, ·	b, ·
<i>middle</i>	c, ·	d, ·
<i>down</i>	e, ·	f, ·

Figure 7. Material payoffs for best response function in a 3 × 2 game in the proof of Theorem 4.

3.4. A Game with Nonlinear Utility Functions

The results of the former section heavily depend on the fact that we restricted ourselves to linear μ - σ utility functions. If we consider other utility functions it might well be that an equilibrium in mixed strategies exists, although the restrictive condition (6) is not met. In order to show this result, we consider the following utility function:

$$V(\mu, \sigma^2) = -\frac{1}{\mu} - \frac{1}{2}\sigma^2.$$

Its indifference curves are convex and monotone functions in the μ - σ^2 diagram.

We now look at a game where the material payoffs are given by Figure 8. We can now show that (0.5, 0.5) is a mixed equilibrium of the game. Notice that in classical game theory (where utilities are given by Figure 8) an equilibrium would be given by (0.75, 0.25).

	<i>left</i>	<i>right</i>
<i>up</i>	3,1	0,0
<i>down</i>	0,0	1,3

Figure 8. Material payoffs of a game with nonlinear utility.

Assume that the row player chooses $q = 0.5$. Then the utility of the column player is given by

$$V(\mu(p), \sigma^2(p)) = -\frac{1}{8} - \frac{3}{2}p + \frac{1}{2}p^2 - \frac{1}{\frac{1}{2} + p}$$

This function has a maximum of $p = \frac{1}{2}$ in $[0, 1]$. Hence, the best response is a mixed strategy. With the same reasoning we can show that $q = \frac{1}{2}$ is the best response to the own player's strategy $p = \frac{1}{2}$.

4. Conclusions

We applied μ - σ utility to game theory. Using monetary games we discussed how equilibria predictions—in particular with respect to the existence of mixed equilibria in static games—change if behavior can be described by preferences depending on the mean and variance of random payoffs. This is an alternative to models of generalized expected utility which relax the assumption of linearity in probabilities, which is the basis of von

Neuman and Morgenstern's expected utility model. While generalized expected utility models still maintain the assumption that terminal utilities are independent of the way the respective endpoint is reached, μ - σ theory allows us to capture endogenous uncertainty caused by mixed strategies of players. In the case of the 2×2 games we were able to show that mixed strategy equilibria do survive in a μ - σ game under a set of additional restrictions. Thus, the set of mixed equilibria in a μ - σ game is a subset of the mixed equilibria of the equivalent game where the monetary payoffs are interpreted as utility payoffs.

Our analysis is based on the interpretation of mixed strategies as randomization by players over their actions, and not as the beliefs about actions of the other player or of the composition of a population of other players playing pure strategies from which the other player is randomly drawn.¹¹ In this case mixed equilibria only survive in μ - σ games if there is a substantial gain from randomizing, for example, because allowing the other player to predict one's behavior comes at a first order cost effect, as in the case of zero sum games.

We believe this analysis can help to capture the experimentally observed aversion against mixing by players. While the μ - σ model is a very specific abstraction and somewhat arbitrary, its prominence in finance and its capability to capture uncertainty endogenous to the play of the game, made it for us a worthwhile starting point to reconsider equilibria when one departs from using utility function that can be characterized by Choquet integrals.

Author Contributions: Methodology, writing, and formal analysis: Both authors contributed equally. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NJ, USA, 1947.
2. Allais, M. Le Comportement de l'Homme Rationnel devant le Risque. *Econometrica* **1953**, *21*, 503–546. [CrossRef]
3. Ellsberg, D. Risk, Ambiguity, and the Savage Axioms. *Q. J. Econ.* **1963**, *75*, 643–669. [CrossRef]
4. Quiggin, J. A Theory of Anticipated Utility. *J. Econ. Behav. Organ.* **1982**, *3*, 225–243. [CrossRef]
5. Yaari, M.E. The Dual Theory of Choice under Risk. *Econometrica* **1987**, *55*, 95–115.
6. Segal, U. The Ellsberg Paradox and Risk Aversion: An Anticipated utility Approach. *Int. Econ. Rev.* **1987**, *28*, 175–202.
7. Cochrane, J.H. *Asset Pricing*, 2nd ed.; Princeton University Press: Princeton, NJ, USA, 2005.
8. Fama, G.; French, K. The cross section of expected returns. *J. Financ.* **1992**, *47*, 427–465. [CrossRef]
9. Markowitz, H.M. Portfolio Selection. *J. Financ.* **1952**, *7*, 77–91.
10. Lajeri-Chaherli, F.; Nielsen, L.-T. Parametric characterizations of risk aversion and prudence. *Econ. Theory* **2000**, *15*, 469–476.
11. Lajeri-Chaherli, F. More on Propriety: The Case of Mean-Variance Preferences. *Geneva Risk Insur. Rev.* **2002**, *27*, 49–60.
12. Meyer, J. Two Moment Decision Models and Expected Utility Maximization. *Am. Econ. Rev.* **1987**, *77*, 421–430.
13. Crawford, V.P. Equilibrium without Independence. *J. Econ. Theory* **1990**, *50*, 127–154. [CrossRef]
14. Chen, H.-C.; Neilson, W.W. Pure-strategy equilibria with non-expected utility players. *Theory Decis.* **1999**, *46*, 19–209. [CrossRef]
15. Rabin, M. Incorporating Fairness Into Game Theory and Economics. *Am. Econ. Rev.* **1993**, *83*, 1281–1302.
16. Battigalli, P.; Dufwenberg, M. Dynamic Psychological Games. *J. Econ. Theory* **2009**, *144*, 1–35. [CrossRef]
17. Choquet, G. Theory of capacities. *Ann. l'Inst. Fourier* **1953**, *5*, 131–295. [CrossRef]
18. Löffler, A. Variance aversion implies μ - σ -criterion. *J. Econ. Theory* **1996**, *69*, 532–539. [CrossRef]
19. Battigalli, P.; Samuelson, L.; Van Huyck, J. Optimization Incentives and Coordination Failure in Laboratory Stag Hunt Games. *Econometrica* **2001**, *69*, 749–764.
20. Al-Ubaydli, O.; Jones, G.; Weel, J. Patience, cognitive skill, and coordination in the repeated stag hunt. *J. Neurosci. Psychol. Econ.* **2013**, *6*, 71–96. [CrossRef]
21. Agranov, M.; Healy, P.J.; Nielsen, K. "Non-Random Randomization", Working Paper. 2020. Available online: <https://ssrn.com/abstract=3544929> (accessed on 11 January 2021).
22. Raiffa, H. Risk, Ambiguity, and the Savage Axioms: Comment. *Q. J. Econ.* **1961**, *75*, 690–694. [CrossRef]

¹¹ With the population interpretation, selection criteria, particularly payoff dominance ([24]), become important. Payoff dominance does not select mixed equilibria in coordination games, as these minimize the expected payoff, and it is difficult to argue why a population's composition should yield this result. μ - σ theory applied to game theory gives us in the case of coordination games a similar prediction.

-
23. Eichberger, J.; Kelsey, D. Uncertainty Aversion and Preference for Randomization. *J. Econ. Theory* **1996**, *71*, 31–43. [[CrossRef](#)]
 24. Harsanyi, J.C.; Selten, R. *A General Theory of Equilibrium Selection in Games*; MIT Press: Cambridge, MA, USA, 1988.
 25. Dekel, E.; Safra, Z.; Segal, U. Existence and dynamic consistency of Nash equilibrium with non-expected utility preferences. *J. Econ. Theory* **1991**, *55*, 229–246. [[CrossRef](#)]
 26. Geanakoplos, J.; Pearce, D.; Stacchetti, E. Psychological Games and Sequential Rationality. *Games Econ. Behav.* **1989**, *1*, 60–79. [[CrossRef](#)]
 27. Nielsen, L.T. Portfolio selection in the mean-variance model: A note. *J. Financ.* **1987**, *42*, 1371–1376. [[CrossRef](#)]