Journal of Choice Modelling 39 (2021) 100282

Contents lists available at ScienceDirect

Journal of Choice Modelling

journal homepage: www.elsevier.com/locate/jocm

Predicting strategic medical choices: An application of a quantal response equilibrium choice model

Ge Ge [a,1,*], Geir Godager [a,b,1]

[a] Department of Health Management and Health Economics, University of Oslo, P.O. Box 1089 Blindern, 0317 Oslo, Norway
[b] Health Services Research Unit, Akershus University Hospital, Sykehusveien 25, 1478 Nordbyhagen, Norway

ARTICLE INFO

JEL classification:
C25
C57
C70
C92
D43
I11

Keywords:
Behavioral game theory
Bounded rationality
Prediction
Quantal response equilibrium
Discrete choice modeling

ABSTRACT

Quantal response equilibrium choice (QREC) models are structural behavioral models that account for bounded rationality and strategic interactions in analyses of games where each player's payoff is a vector. We revisit the question of how market competition affects pro-social behavior and fit a QREC model to data from an incentivized laboratory experiment, where participants make decisions on medical treatments for abstract patients in *monopoly*, *duopoly*, and *quadropoly* games. Our results demonstrate that competition can cause substantial behavioral responses without any changes in pro-social preferences if one allows for the possibility that competition influences the degree of randomness in decision making.

We find that a QREC model with fixed preference parameters provides precise out-of-sample predictions of behavior in games with vector payoffs. A Monte Carlo study is performed to show that the two-step estimator is accurate.

1. Introduction

In their seminal contribution, McKelvey and Palfrey (1995) proved that a quantal response equilibrium (QRE) always exists for finite games. A QRE can be described as a statistical version of a Nash equilibrium (Camerer, 2011), where sub-optimal choice alternatives have non-zero choice probabilities. By allowing for sub-optimal strategic choices, the QRE introduces a weaker assumption on human rationality to game theory. Furthermore, McKelvey and Palfrey (1995) bridged the gap between Behavioral Game Theory and Choice Modeling by showing that behavior by individuals who maximize a linear combination of expected utility and noise is a game theoretic equilibrium. The objective in the QRE model by McKelvey and Palfrey (1995) can be written as:

$$U_j = \lambda V_j + \varepsilon_j \quad , \tag{1}$$

where expected utility of choosing pure strategy (alternative) j is denoted by V_j, and ε_j are the noise terms. Assuming ε_j to be independent and extreme value distributed (type I) implies a logistic quantal response equilibrium (LQRE) model. The λ-parameter is the scale parameter of a multinomial logit (MNL) model, denoted by σ in Fiebig et al. (2010). This scale parameter is typically referred to as the "rationality parameter" in the behavioral game theory literature inspired by McKelvey and Palfrey (1995).

Twenty-five years after McKelvey and Palfrey (1995), most applications are still focusing on strategic scenarios where the choice alternatives are characterized by a single attribute. Generalizations to study strategic agents caring about multiple attributes are

* Corresponding author.
E-mail addresses: gege@medisin.uio.no (G. Ge), geir.godager@medisin.uio.no (G. Godager).

[1] Ge Ge and Geir Godager contributed equally.

https://doi.org/10.1016/j.jocm.2021.100282
Received 21 April 2020; Received in revised form 22 January 2021; Accepted 11 February 2021
Available online 27 February 2021
1755-5345/

straight forward and familiar to choice modelers. Yet, there are few choice modelers applying quantal response equilibrium choice (QREC) models, i.e. QRE models where decision-makers have preferences for multiple attributes. One of the aims of this paper is to inspire the use of QREC models in applied research. We show how ready-made software can be applied to obtain a two-step QREC estimator, and conduct Monte Carlo simulations to provide documentation on its precision.[2] Not only is this two-step QREC estimator a convenient approach, the results from our Monte Carlo simulations show that it is also accurate even with a moderate sample size.

We conduct an incentivized lab experiment to study strategic behavior under various levels of competition. The experiment is framed in a medical setting where the individuals play the role of physician. In the game, the subjects make treatment choices that generate profits for themselves as well as positive health benefits for the patients. Three levels of competition are considered: monopoly, duopoly, and quadropoly. In the competitive scenarios, the subjects' joint choices of treatment alternatives determine the subjects' profits as well as the health benefits for patients in the market.

The results show that the substantial difference in observed behavior under different competition intensities can be attributed to a change in the scale parameter. In particular, the choices become less random as the competition intensifies. This finding provides useful nuance to the recent literature on whether markets erode social responsibility (Falk and Szech, 2013; Bartling et al., 2015; Kirchler et al., 2015; Bartling et al., 2021) and contributes to the small empirical literature on physician competition (Gaynor and Vogt, 2000; Gaynor and Town, 2011).

The paper proceeds as follows: In Section 2, we give a brief overview of recent literature on the scale in choice models and behavioral game theory. In Section 3, we apply a QREC model to a physician oligopoly and describe the two-step estimation procedure. In Section 4, we give a brief description of the experimental protocol. Descriptive statistics and results from estimation and Monte Carlo simulations are presented in Section 5. Finally, we discuss and conclude in Section 6.

2. Related literature

2.1. Stochastic choice

As described by McFadden et al. (1999), random utility models allow for weaker forms of human rationality than the strong rationality assumption typically employed in textbook economics, and theoretical contributions have shown that logit specifications are highly non-restrictive.[3] The MNL model in (1) is highly versatile and can be motivated and deduced in many different ways:

- According to Thurstone (1927a,b), randomness in behavior stems from individuals' utilities varying from moment to moment in a stochastic manner.
- A different way of motivating (1) is to regard individuals as incapable of making perfectly rational choices. In the model deduction by Luce (1959), Tversky (1972), and McKelvey and Palfrey (1995), the individual's (expected) utility is assumed deterministic, while randomness occurs in implementation of choice. Luce describes a kind of perception error (McFadden et al., 1999) in motivating randomness. He proposes that individuals are unable to discriminate perfectly between utility levels of available alternatives.[4]
- Swait and Marley (2013) and Wallin et al. (2018) show that (1) follows by implication from the optimizing behavior of a decision-maker balancing the competing goals of achieving high utility and product variation.
- Erlander (1998) shows that (1) can be motivated by the implication from "the efficiency assumption" that samples with higher total observable utility are more probable.
- The MNL model has often been presented as a practical econometric specification where the noise term is introduced to account for variables that are unobservable to the analyst (McFadden, 1974).

The fact that there are many ways to motivate and deduce (1) does not mean that a given way of motivating the MNL model suits every purpose. Introducing the noise term as *unobservables* appear to be less convincing when analyzing data collected in a controlled lab experiment. Using the motivation by Thurstone (1927a,b) or Swait and Marley (2013) appears to be less plausible for models of (errors in) medical decision making (Mackowiak, 2020), a pilot's choice of applying flaps for takeoff (Loukopoulos et al., 2009), or losses by Grandmasters in chess.

The bounded rationality perspective is the motivation applied by McKelvey and Palfrey (1995), and they interpret the scale parameter as individuals' *degree of rationality*. This interpretation can be criticized, as one might argue that *rationality* would have to be very narrowly defined to be represented by one single parameter. Furthermore, for investment decisions and many other examples, randomizing can be a rational strategic choice. Regardless of motivation, the scale parameter is a measure of the *degree of randomness in behavior*. We take this broad perspective in this paper.

[2] User written programs for estimating generalized multinomial logit models (gmnl) are available for Stata (Gu et al., 2013) and R (Sarrias and Daziano, 2017).

[3] See Dagsvik (2016) and the references therein.

[4] The difference in motivation between Luce and Thurstone becomes superficial in practical applications of standard choice models. McFadden (1981) shows that the two types of probabilistic choice models are equivalent in many cases.

2.2. Scale in choice models

Controlled lab experiments provide favorable conditions for identifying the scale parameter in choice models. Experiments performed in a controlled environment facilitate implementation of ceteris paribus changes in the choice context and enable the researcher to confront each decision-maker with the exact same set of choice scenarios. A context-dependent scale parameter is identified under the assumption that preference parameters are fixed and independent of the context.

The scale in choice models has generated much confusion (Swait and Louviere, 1993; Louviere and Eagle, 2006; Hess and Rose, 2012; Hess and Train, 2017). Empirical contributions have often ignored or avoided addressing questions about how the scale might differ systematically between individuals or between choice occasions for the same individual. An important practical reason is that many available data sets might not contain sufficient information to address questions regarding, for example, scale differences caused by systematical contextual differences between choice occasions. While data collected in the field may record context differences, such as whether commuters on public transport were interviewed in rush hour, it would rarely be the case that decision-makers are randomly assigned to a choice context. As a result of unobserved selection mechanisms, it is highly plausible that individuals sampled at different choice contexts differ on both preference parameters and scale.

Louviere and Eagle (2006) argue that scale is highly unlikely to be constant, as the impact of noise on choices can vary over conditions, context circumstances, or situations, as well as between decision-makers. Furthermore, as illustrated with examples by Louviere and Eagle (2006), if scales vary across decision-makers, decision-makers will seem to be heterogeneous in preferences even if they differ only in scale. For the same reason, differences in scale across choice occasions can make individuals' preferences appear to be context-dependent even when preferences are stable.

As discussed by Louviere et al. (1999), whether preferences *are* context-dependent or *appear* context-dependent as a result of ignoring between-context differences in scale, is an important question for choice modelers. Evidence of stable preferences provides support for the validity of results from stated preference experiments when aiming to predict market behavior. The question of whether preferences are stable is equally important for experimental economists aiming to predict behavior in a future experiment. One cannot generally disregard the possibility that preferences are context dependent. Firms and governments allocate resources to commercial marketing services, aiming to shape consumer preferences by repeated exposure to certain stimuli. One may argue that a lab experiment of the type applied in this study will rarely be life-changing for the participants. We, therefore, assume that the individuals' valuation of health benefits for patients is predetermined, and that the participants' values are unaffected by taking part in a laboratory experiment where they simply choose an abstract treatment alternative on 24 different occasions.

2.3. Empirical game theory

The Nash equilibrium (NE) (1950) is useful in providing theoretical predictions in single-criterion games and has undoubtedly been pivotal to the development of game theory and the empirical analysis of strategic behavior. The fact that the NE fits poorly with observed behavior in many cases is much discussed in the literature on empirical game theory, with an example being the enlightening contribution by Goeree and Holt (2001). For decades, the poor fit of the NE assumption inspired explorations of new equilibrium concepts that could explain the observed behavior (Harsanyi, 1973; Harsanyi et al., 1988; Ma and Manove, 1993).

Strategic economic agents often care about more than a single criterion when making economic choices. Firms are concerned about both long-run and short-run profits, whereas politicians deciding on foreign trade policies might be concerned about several aspects of policy impacts, such as effects on the environment, unemployment, budget deficits as well as the likelihood of re-election. Medical doctors who compete for patients, as in the framing of the experiment studied in this paper, are likely concerned about profits as well as health effects for patients when choosing medical treatments. A game with vector payoff is often referred to as a multi-criteria game, or a multi-payoff game (Shapley and Rigby, 1959; Zeleny, 1975; Voorneveld et al., 2000), and would perhaps be translated to a "multiple attribute game" in choice modeling terminology. If the players' valuation of payoff elements is known to the researcher, the multi-criteria game can be scalarized into a standard single-criterion game, which can be solved for NE by using standard approaches.

Despite the obvious relevance of multi-criteria games for analyzing strategic behavior when the payoff is a vector, its literature is small compared to that of the single-criterion game theory. A plausible explanation is the lack of useful equilibrium concepts for multi-criteria games. Although equilibrium concepts, such as, the *Pareto equilibrium* (Shapley and Rigby, 1959; Voorneveld et al., 1999) and *ideal equilibrium* (Voorneveld et al., 2000) have been developed for solving such games, there are clear limitations when individuals' preferences are latent. Shapley and Rigby (1959) show that narrowing down the set of plausible actions in multi-criteria games is difficult when the individuals' valuation of the elements in the vector payoff is unknown. Therefore, the *Pareto equilibrium* discussed by Shapley and Rigby (1959) is often not useful, as many plausible actions might remain after removing dominated strategies. Furthermore, the number of pure strategy NEs and whether a given set of actions by the players constitute pure strategy NE can depend on unknown preference weights.

In their seminal contribution, McKelvey and Palfrey (1995) proved the existence of a quantal response equilibrium. As described by Jessie and Saari (2016), one of the contributions of McKelvey and Palfrey (1995) is linking behavioral game theory to the choice modeling paradigm. Choices in QRE models are assumed to be the result of individuals maximizing a linear combination of expected utility and noise as in (1).[5] The unknown value of λ has motivated much criticism of the QRE assumption. While the QRE assumption

[5] With a structure similar to Swait and Marley (2013), one could regard a QRE as an equilibrium of a multi-criteria game where the players' weighing of the two criteria, a scalar payoff and noise, is unknown. In this text, however, the term multi-criteria refers to games where payoffs are vectors of observable attributes.

Table 1
The matrix of payoff vectors to the row player in a symmetric duopoly game with two pure strategies.

		The opponent	
		L	H
Row	L	$\Pi_{L\|L}, B_{L\|L}$	$\Pi_{L\|H}, B_{L\|H}$
player	H	$\Pi_{H\|L}, B_{H\|L}$	$\Pi_{H\|H}, B_{H\|H}$

is shown to fit empirical data in many applications, it has been criticized for being difficult to falsify. As noted by Haile et al. (2008), an estimated scale parameter in a particular game might not be useful in predicting behavior in other games or by other players. For this reason, some authors, such as McCubbins et al. (2013), even argue that QRE *"...is of limited use in explaining human behavior across even a small range of similar decisions."* However, other researchers, such as Anderson et al. (2001), Goeree et al. (2005), Goeree and Holt (2005), Goeree et al. (2010), Matějka and McKay (2015) and Wright and Leyton-Brown (2017), are more optimistic in their view on the role of QRE in future empirical research on strategic decision making. Wright and Leyton-Brown (2017) suggest an agenda for acquiring more knowledge on the usefulness of the QRE assumption when aiming to provide ex-ante predictions of strategic behavior. They propose to study carefully the out-of-sample predictions of models assuming QRE. Using an estimated scale parameter from one game to predict behavior in a different game, we may acquire new knowledge on the usefulness of the QRE assumption.

3. A quantal response equilibrium choice model

3.1. The basic 2 × 2 game for two players

For clarity and ease of exposition, we start with the most basic example of a strategic scenario: A symmetric two-player game where two identical agents have an identical binary choice set. Our game differs from more commonly studied games where the players' payoff is a scalar. In our game, the payoff is a vector, and the values of its elements are determined by the players' joint choices. The two players comprise a duopoly of physicians, who can choose between either Low or $High$ treatment intensities for their patients. The two physicians' joint choices of treatment plans determine the vector of profit and patient health benefits that each of the two players receives. We assume that physicians value both profit, Π; and patient's benefit, B.

The game is described in a normal form representation with a row player and a column player in Table 1. $\Pi_{r|c}$ and $B_{r|c}$ denote the payoff to the row player when he chooses row r and the column player chooses column c, $r, c = L, H$. We study symmetric games in this paper, so reporting the payoffs to the row player is sufficient for describing the game. In a symmetric two-player game, the payoff to the row player when he chooses treatment L and the column player chooses treatment H is identical to the payoff received by the column player in the mirrored situation where the column player chooses treatment L and the row player chooses treatment H. We assign the first-person perspective to the row player and let all payoff vectors represent payoffs to the row player. In this way, we economize on notation and highlight the common ground of strategic decision making and choices under uncertainty.

The utility for the row player of choosing, for example, L is uncertain, as it depends on whether the opponent chooses L or H. Conditional on the column player choosing c, the row player's utility from choosing row r, modeled as a simple linear-in-parameter preference function, can be expressed as:

$$v_{r|c} = \beta_\Pi \Pi_{r|c} + \beta_B B_{r|c} \quad . \tag{2}$$

Assuming observed choices to be the result of individuals maximizing a linear combination of *(objective) expected utility* and noise involves only a small and natural augmentation of the basic choice model in (1).[6] Let P_L and P_H denote the unobserved probabilities that the opponent chooses alternative L or H, respectively. Given the observed payoff matrix in Table 1, the row player's expected utility from choosing row r can be expressed as:

$$V_r(\mathbf{P}) = P_L v_{r|L} + P_H v_{r|H}. \tag{3}$$

Inserting (2) into (3), and rearranging terms, we get:

$$V_r(\mathbf{P}) = P_L(\beta_\Pi \Pi_{r|L} + \beta_B B_{r|L}) + P_H(\beta_\Pi \Pi_{r|H} + \beta_B B_{r|H}) \tag{4a}$$
$$= \beta_\Pi(P_L \Pi_{r|L} + P_H \Pi_{r|H}) + \beta_B(P_L B_{r|L} + P_H B_{r|H}) \tag{4b}$$
$$= \beta_\Pi E\Pi_r + \beta_B EB_r \tag{4c}$$

In (4c), we have introduced notation for the expected attribute vector of r: $E\Pi_r = P_L \Pi_{r|L} + P_H \Pi_{r|H}$ and $EB_r = P_L B_{r|L} + P_H B_{r|H}$. Eq. (4) highlights that expected utility is a function of the unobserved multinomial probability vector $\mathbf{P} = (P_L, P_H)$.

[6] There is a growing body of literature on *Random Expected Utility Models*. See, for example, Gul and Pesendorfer (2006), Blavatskyy (2007), Dagsvik (2008, 2015) and Ke (2018).

Inserting $V_r(\mathbf{P})$ for V_r in (1) and assuming ε_r to be type 1 extreme value distributed, the probability that the row player chooses L becomes:

$$P_L = \frac{e^{\lambda V_L(\mathbf{P})}}{e^{\lambda V_L(\mathbf{P})} + e^{\lambda V_H(\mathbf{P})}} \quad , \tag{5}$$

where λ is the scale parameter. As $\lambda \longrightarrow 0$, the behavior becomes purely random, and the player plays each pure strategy with equal probability. As $\lambda \longrightarrow \infty$, the player makes no errors, and the behavior becomes deterministic. In the limit, the players are perfectly responsive to the differences in expected utility across alternatives, and the behavior converges to a NE.

Data generated from the game in Table 1 differ from typical choice data in that the expected attribute vector in (4c) is unobservable. This is a key feature that distinguishes (5) from the standard scaled logit model. An intuitive two-step estimation procedure can be applied to estimate (β, λ). The first step is to replace unobserved probabilities $\mathbf{P}$ on the right-hand side of (4c) with observed relative choice frequencies, $\mathbf{f}$, to obtain an estimate of the expected attribute vector. The second step is to estimate (β, λ) by means of maximum likelihood, treating the estimated attribute vector as if it was an observed attribute vector. In the following we refer to this as a two-step QREC estimator. An important note is that a choice modeler tasked with analyzing the data from this game might likely proceed with the two-step procedure without being knowledgeable of McKelvey and Palfrey's contribution of QRE.[7]

A key feature of (5) from a behavioral game theory perspective is that $\mathbf{P}$ is a function of itself, as the left-hand side of (5) is one of the elements of $\mathbf{P}$ on the right-hand side. McKelvey and Palfrey (1995) proved that a QRE always exists for a finite game, which means that a fixed-point solution of (5) is guaranteed to exist. While solving the fixed point of (5) or having knowledge of the existence of QRE is not necessary for obtaining the two-step estimator, it should be reassuring for applied choice modelers that a QRE always exists in a game with well-defined choice sets.

In the following subsection, we introduce functional form and generalize the QREC model to more than two players and J pure strategies.

3.2. Generalized model

There are three different market settings denoted by t, t =[monopoly, duopoly, quadropoly], and eight independent games are played within each market setting. We present the payoff matrices of the 24 games in Appendix A and elaborate one example in the next section. A complete description of the formulas for computing payoff matrices based on the experimental parameters can be found in Ge and Godager (2021).

Consider a physician selecting one treatment alternative j from a finite set $C = 1, 2, \ldots, J$. The choice set comprises J mutually exclusive alternatives within a game. To simplify notation, in this subsection, we suppress the index for each game within a market setting. We assume that the physicians' preference parameters are homogeneous and fixed over games. While assuming a representative individual model appears restrictive after decades of applications of mixed logit models, a parsimonious model specification is suitable for our purposes.[8]

The payoff players receive is a vector consisting of two elements: physician profit, Π; and patient benefit, B. Similar to Goeree et al. (2002), we assume that a healthy patient population in the market is a shared good, and a physician's valuation of the patient benefit in the market is independent of which physician provides treatment to which patient. In other words, B here is the total benefit of all the patients in the market. We represent preferences by a linear-in-parameter function of the two payoff elements. We consider two alternative specifications: a quadratic function and a Cobb–Douglas function in a log-linear form. While the quadratic specification has obvious drawbacks, such as saturation, it is convenient and simple. The log-linear Cobb–Douglas arguably has a more solid axiomatic foundation following the proof by Dagsvik (2018) but offers less flexibility than the quadratic specification. The non-linear utility specification allows us to identify the curvature of the utility function (see, for example, Van Der Pol et al. (2010), Kolstad (2011), van der Pol et al. (2014), Holte et al. (2016), Wang et al. (2020)). By Taylor's theorem, further expanding the polynomial in the two specifications would provide better approximations. However, such improvements in functional forms are costly, as more and richer data are required to quantify additional parameters. In addition, larger samples and additional parameters also raise computational costs. Hence, a quadratic form and log-linear Cobb–Douglas, rather than a more general translog, is a convenient choice.

The payoff elements of treatment alternative j in monopoly is denoted by Π_{jt} and B_{jt}. In duopoly, $\Pi_{jt|x}$ and $B_{jt|x}$ are the payoff elements of choosing j given that the opponent's choice is x. In quadropoly, $\Pi_{jt|xyz}$ and $B_{jt|xyz}$ are the payoff elements of choosing j given that the combination of choices by opponents are xyz. The conditional utility of a physician choosing alternative j given the opponent(s)' choice(s) is denoted by v_{jt}, $v_{jt|x}$, and $v_{jt|xyz}$ in monopoly, duopoly, and quadropoly, respectively. Preference parameters are denoted by the vector β. The conditional utility for the two functional forms in, for example, quadropoly can be written as:

[7] This was the case for the authors of this paper.

[8] While generalizing the model to account for the unobserved heterogeneity in λ or preference parameters is an obvious extension, specifying a representative individual model is a convenient choice as it substantially reduces the burden of conducting the Monte Carlo simulations of repeated equilibrium game play. Our parsimonious specification can also be defended by arguing that preferences for the attributes in the experiment (money and health benefit) are likely to be less heterogeneous than preferences for, for example, household appliances.

$$\textbf{Quadratic utility}: \quad v_{jt|xyz} = \beta_{\Pi}\Pi_{jt|xyz} + \beta_B B_{jt|xyz} + \beta_{\Pi\Pi}\Pi^2_{jt|xyz} + \beta_{BB}B^2_{jt|xyz} + \beta_{\Pi B}\Pi_{jt|xyz}B_{jt|xyz}, \tag{6}$$

$$\textbf{Cobb–Douglas utility}: \quad v_{jt|xyz} = \underline{U} + \beta_{\Pi} ln\Pi_{jt|xyz} + \beta_B ln B_{jt|xyz}, \tag{7}$$

where $j, x, y, z \in C$ and t =[monopoly, duopoly, quadropoly], and $\underline{U}$ is a reference utility.

The opponent(s)' action(s) are unknown ex ante. The probability that the opponent chooses alternative x in duopoly is denoted by P_{xt}, and the probabilities of the three opponents choosing x, y, and z in quadropoly is denoted by P_{xt}, P_{yt} and P_{zt}, respectively. When a physician chooses alternative j, his expected utility, $V_{jt}(\mathbf{P}_t)$, is given by:

$$\text{Monopoly: } V_{jt}(\mathbf{P}_t) = v_{jt} \; ; \tag{8a}$$

$$\text{Duopoly: } V_{jt}(\mathbf{P}_t) = \sum_{x\in C} P_{xt} v_{jt|x} \; ; \tag{8b}$$

$$\text{Quadropoly: } V_{jt}(\mathbf{P}_t) = \sum_{x\in C}\sum_{y\in C}\sum_{z\in C} P_{xt}P_{yt}P_{zt} v_{jt|xyz} \; , \tag{8c}$$

where $\mathbf{P}_t = (P_{1t}, P_{2t}, \dots, P_{Jt})$ denotes the vector of the opponent(s)' choice probabilities for all alternatives in market setting t. We generalize (1) by assuming that the physician maximizes a linear combination of the expected utility from choosing j and an error term, ε_{jt}, in each game:

$$\lambda_t V_{jt}(\mathbf{P}_t) + \varepsilon_{jt} \; , \tag{9}$$

In our symmetric games with homogeneous players, the probability of choosing j in market setting t is the same for each player and given by:

$$P_{jt} = \frac{e^{\lambda_t V_{jt}(\mathbf{P}_t)}}{\sum_{r\in C} e^{\lambda_t V_{rt}(\mathbf{P}_t)}}, \tag{10}$$

One cannot identify both the scale parameter and the preference parameters (Train, 2009; Hess and Rose, 2012). In the quadratic utility specification, we normalize $\beta_{\Pi} + \beta_{\Pi\Pi} = 1$ in order to identify λ in all three markets.[9] In the Cobb–Douglas specification, we apply the normalization $\beta_{\Pi} + \beta_B = 1$.[10] We handle zeros in the Cobb–Douglas specification by introducing a dummy equal to one whenever $\Pi = 0$ or $B = 0$. This dummy captures the reference utility $\underline{U}$ in (7), and it is switched on whenever we replace $ln(0)$ by 0.[11]

3.3. Estimation

We now let $g = 1, 2, \dots, 8$ index the eight patient games in each market setting, and $n = 1, 2, \dots, 136$ index the individuals. The log-likelihood function given our QREC model specification can be written as:

$$LL(\lambda_t, \beta)|_{\mathbf{P}_{gt}} = \sum_n \sum_j \sum_g \sum_t y_{njgt} ln\left(\frac{e^{\lambda_t V_{jgt}(\mathbf{P}_{gt},\beta)}}{\sum_{r\in C} e^{\lambda_t V_{rgt}(\mathbf{P}_{gt},\beta)}}\right), \tag{11}$$

where $y_{njgt} = 1$ if physician n chose j in patient game g in market setting t, and zero otherwise. If $\mathbf{P}_{gt}$ is known, preference estimates from maximizing equation (11) can be acquired directly.

In this paper, we use the two-step procedure, where $\mathbf{P}_{gt}$ is replaced by relative frequencies, $\mathbf{f}_{gt}$, and document the accuracy by means of Monte Carlo simulations.[12] With $\mathbf{P}_{gt}$ replaced by $\mathbf{f}_{gt}$, the likelihood function can be written:

$$LL(\lambda_t, \beta)|_{\mathbf{f}_{gt}} = \sum_n \sum_j \sum_g \sum_t y_{njgt} ln\left(\frac{e^{\lambda_t V_{jgt}(\mathbf{f}_{gt},\beta)}}{\sum_{r\in C} e^{\lambda_t V_{rt}(\mathbf{f}_{gt},\beta)}}\right). \tag{12}$$

Since $\mathbf{f}_{gt}$ is a consistent estimator of $\mathbf{P}_{gt}$, (12) converges in probability to Eq. (11) as the number of individuals increases towards infinity. It follows that the estimators $\hat{\lambda}_t$ and $\hat{\beta}$, which maximize Eq. (12), are consistent. They are estimated by the `gmnl` module in Stata 16 (Gu et al., 2013). Note that different from the full information likelihood procedure described by Moffatt (2015), computing the fixed point, assuming subjects play a LQRE, is not necessary for using the two-step procedure.

9 In this way, the estimates are measured in the marginal utility of the first dollar of profit when the health benefit is zero.

10 See, e.g., Swait and Marley (2013) for an example of constraining the sum of the coefficients.

11 A similar procedure is described by Battese (1997).

12 Although the two-step procedure is frequently used (see, e.g., Bajari and Hortacsu, 2005), we are unaware of any studies that document the accuracy of the two-step estimator.

Table 2
Payoff matrix, patient game 1 in duopoly.

	The opponent										
	0	1	2	3	4	5	6	7	8	9	10
Player n	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B	Π, B
0	500,0	270,73	120,176	50,285	20,392	10,495	0,600	0,700	0,800	0,900	0,1000
1	723,73	495,100	267,173	119,276	50,385	20,492	10,595	0,700	0,800	0,900	0,1000
2	845,176	701,173	480,200	259,273	115,376	48,485	19,592	10,695	0,800	0,900	0,1000
3	865,285	801,276	664,273	455,300	246,373	109,476	46,585	18,692	9,795	0,900	0,1000
4	823,392	798,385	739,376	613,373	420,400	227,473	101,576	42,685	17,792	8,895	0,1000
5	743,495	735,492	713,485	660,476	548,473	375,500	203,573	90,676	38,785	15,892	8,995
6	640,600	634,595	627,592	608,585	563,576	467,573	320,600	173,673	77,776	32,885	13,992
7	510,700	510,700	505,695	500,692	485,685	449,676	372,673	255,700	138,773	61,876	26,985
8	360,800	360,800	360,800	356,795	353,792	342,785	317,776	263,773	180,800	97,873	43,976
9	190,900	190,900	190,900	190,900	188,895	186,892	181,885	167,876	139,873	95,900	51,973
10	0,1000	0,1000	0,1000	0,1000	0,1000	0,995	0,992	0,985	0,976	0,973	0,1000

Payoffs are measured in Taler (100 Taler = 1 Euro).

4. Experimental data

In this section, we briefly describe the experimental design and show an example of the payoff matrices with illustration. A complete description of the experimental protocol and the formulas for computing payoff matrices based on the experimental parameters can be found in Ge and Godager (2021).

4.1. Experiment

The experiment has a medical framing. Participants are instructed to play the role of a physician and choose one of eleven different medical treatments in eight patient games in each of three markets with various levels of competition: monopoly, duopoly, and quadropoly. In the instructions for the participants, we use the term "patient type" rather than patient game (see Appendix B for instructions for participants). The trade-offs between profit and patient benefit depend on the payoff matrix, which varies across patient games and market settings. The differences in payoffs between market settings for a given patient game are caused by differences in the market shares. In duopoly and quadropoly, the decision-makers' joint treatment choices determine their own profits and the patients' health benefits.[13] The profit and benefit accrued in the laboratory are converted into monetary transfers to the participants and a charity dedicated to providing surgeries for ophthalmic patients. This element of our protocol, which is identical to Hennig-Schmidt et al. (2011), motivates participants' patient-regarding behavior in the laboratory. The experiment was designed to accommodate independent games of complete information with simultaneous moves. It was programmed in z-Tree (Fischbacher, 2007). Funding for the development of experimental design, the programming, as well as the payments to participants and *Christoffel Blindenmission* was provided by the Research Council of Norway. The ethical review and approval of experimental procedure were given by the Norwegian Social Science Data Services (reference No. 43709).

4.2. Payoff matrices in three markets

We now present the payoff matrix for patient game 1 in duopoly in Table 2 as an example, and a complete set of payoff matrices of the experiment is provided in Appendix A. Since the game is symmetric, we only present player n's (the row player) payoffs. Each cell represents a vector of profit and patient benefit that player n receives given the combination of treatments offered by player n and his opponents. If, for instance, player n chooses alternative 3 for patient 1 in duopoly, and the opponent chooses alternative 2, then player n will receive 664 Taler of profit, and 273 Taler will be provided to the patients (100 Taler = 1 Euro). The subjects could inspect the payoff elements in each cell of the matrix by means of a "calculator" when choice combinations by players were inserted (Requate and Waichman, 2011).[14]

5. Results

5.1. Descriptive statistics

We report the observed choice frequencies in Table 3 and the relative frequencies in Fig. 2.[15] We see in Table 3 that 24 out of 136 subjects chose pure strategy 0 for patient game 1 in monopoly, whereas none of the decision-makers chose this alternative for patient game 1 in duopoly and quadropoly. Health benefit to patients is proportional to the numbering of the pure strategies, so that

[13] Interested readers can find details on parameters for patient games and formulas for computing payoff matrices for 24 games in Ge and Godager (2021).
[14] Interested readers may also calculate the payoff matrix using the formulas and game parameters given in Ge and Godager (2021).
[15] Note that Table 3 and the payoff matrices in Appendix A are sufficient for reproducing all empirical results in this paper.

Table 3
Observed frequencies of strategy choice in the 24 games by the 136 subjects.

Market	Patient game	Pure strategy										
		0	1	2	3	4	5	6	7	8	9	10
Monopoly	1	24	12	4	18	14	27	16	15	2	2	2
	2	21	9	9	13	11	18	24	14	10	1	6
	3	23	6	10	9	11	14	10	22	14	6	11
	4	23	5	7	13	11	19	9	29	11	3	6
	5	24	7	7	19	18	15	14	21	5	4	2
	6	23	4	8	21	14	17	11	15	13	4	6
	7	21	6	7	14	7	14	11	13	24	7	12
	8	21	4	14	12	7	16	14	8	20	7	13
Duopoly	1	0	0	1	3	3	12	24	41	26	22	4
	2	0	0	0	2	4	7	12	27	36	26	22
	3	0	0	0	3	4	3	9	18	37	34	28
	4	1	0	3	1	0	4	6	21	30	29	41
	5	1	0	0	1	4	18	22	48	32	7	3
	6	1	0	0	1	4	9	14	33	42	19	13
	7	1	1	0	1	1	1	3	15	18	28	67
	8	0	1	0	0	1	2	7	14	25	28	58
Quadropoly	1	0	0	0	2	2	4	14	31	48	30	5
	2	0	0	0	2	0	3	10	11	42	31	37
	3	1	0	0	0	0	0	3	12	26	41	53
	4	1	0	1	0	0	0	4	9	24	39	58
	5	0	0	0	0	0	10	9	37	58	17	5
	6	0	0	0	1	1	1	8	15	40	43	27
	7	0	0	0	1	0	1	3	2	16	22	91
	8	0	0	0	0	1	1	4	5	14	27	84

Table 4
Results from estimation.

Model	Quadratic utility[a]			Cobb–Douglas utility		
	Estimates	95% C.I.[b]		Estimates	95% C.[b]	
Preference parameters						
β_{Π}	1.05	1.04	1.05	0.69	0.67	0.71
β_{B}	0.59	0.46	0.73	0.31	0.29	0.33
$\beta_{\Pi\Pi}$	−0.05	−0.05	−0.04	n.a.	n.a.	n.a.
β_{BB}	−0.03	−0.04	−0.02	n.a.	n.a.	n.a.
$\beta_{\Pi B}$	−0.05	−0.05	−0.04	n.a.	n.a.	n.a.
$\underline{U}$	n.a.	n.a.	n.a.	1.88	1.77	1.98
λ Parameters						
$Monopoly$	0.54	0.40	0.69	1.80	1.39	2.21
$Duopoly$	2.39	2.13	2.64	2.17	1.98	2.37
$Quadropoly$	4.13	3.76	4.50	2.51	2.30	2.73
Log-Likelihood	−6121.90			−6243.36		
# subjects:	136					
# Games:	24					

[a]Variables in the quadratic specifications were scaled in Euro instead of Taler.
[b]Estimated 95% confidence interval based on standard errors from the maximum likelihood estimation as reported by the software. Parameter normalizations: $\beta_{\Pi} + \beta_{\Pi\Pi} = 1$ in the quadratic specification, and $\beta_{\Pi} + \beta_{B} = 1$ in the Cobb–Douglas specification.

when more decision-makers pick alternatives in the columns further to the right in Table 3, more health benefits are transferred to the charity. Hence, from observing Table 3 and Fig. 2, it is clear that, in this experiment, more competition benefits the patients in the market. This finding is in contrast with Falk and Szech (2013). See Section 6 for further discussion.

5.2. Estimation results

We present the estimated scale and preference parameters in Table 4. While all estimated preference parameters are reported to be significant, the reported standard errors need to be interpreted in a conditional sense, as estimates of $\mathbf{P}_{gt}$ are used for approximating the likelihood function. The estimated preference parameters from the Cobb–Douglas specification are similar in magnitude to those reported by Wang et al. (2020), who fit a scaled logit model with log-linear Cobb–Douglas utility to the data from experiments based on an extended version of the design by Hennig-Schmidt et al. (2011). The confidence interval for the preference parameter in Table 4 overlaps with the confidence interval in Wang et al. (2020).

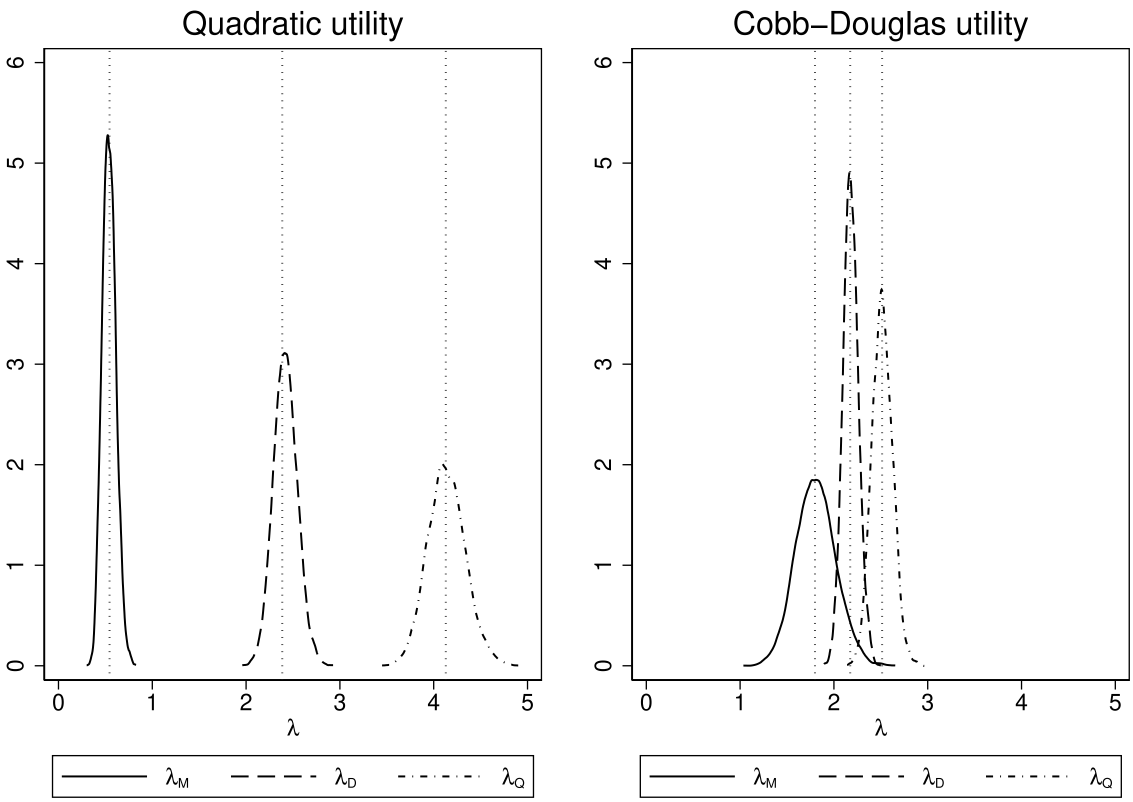

Fig. 1. Distribution of λ estimates from Monte Carlo simulation. Dotted vertical lines indicate the given parameter values.

Assuming the estimated standard errors are accurate, we apply simple Wald tests to test the null hypothesis that the three scale parameters are identical across market settings. We may reject this hypothesis for both quadratic (p-value < 0.0001) and Cobb–Douglas (p-value < 0.0007) specifications. We also perform pairwise tests. For the quadratic specification, we may reject the null hypothesis that scale parameters in monopoly and duopoly are equal (p-value < 0.0001) and the null hypothesis that the scale parameters in duopoly and quadropoly are equal (p-value < 0.0001). For the Cobb–Douglas specification, we cannot reject the null hypothesis that the scale parameters in monopoly and duopoly are equal (p-value $= 0.0856$). However, we may reject the null hypothesis that the scale parameters in duopoly and quadropoly are equal (p-value $= 0.0056$). The interpretation is that higher competition raises the scale parameter. We return to the interpretation of this result in the discussion section.

For the quadratic specification, we observe that the estimates of $\beta_{\Pi\Pi}$ and β_{BB} are negative, implying that the marginal utilities are declining in profit and patient benefit. In addition, the estimate of $\beta_{\Pi\mathrm{B}}$ is also negative, indicating that profit and patient benefit are substitutes for this utility specification.

Conditional on parameter estimates from Table 4 and the 16 payoff matrices where competition is present (Tables A.2–A.9 and Tables A.10–A.17 in Appendix A), we can compute the matrices comprising the scalar utility payoffs in each game (see Appendix C for the computed utility payoff games assuming the quadratic preferences). We see that the numbers of dominated strategies and NE differ substantially between games. For example, the pure strategy Nash equilibrium in patient game 1 in duopoly is unique, but we find many cases where the scalarized games have multiple NEs, such as patient game 7 in duopoly with six different NEs.

5.3. Monte Carlo simulation

We conduct Monte Carlo simulations to assess the precision of the two-step QREC estimator. Let $(\hat{\beta}\ ,\ \hat{\lambda}_t)$ denote our estimated preference and scale parameters, which are presented in Table 4. We apply these parameters as the fixed parameters in the Monte Carlo simulations. From the proof by McKelvey and Palfrey (1995), we know that the preference and scale parameter $(\hat{\beta},\ \hat{\lambda}_t)$ determine a LQRE given the payoff matrix for the occasion indexed gt. We denote this LQRE by $\hat{\mathbf{P}}^*_{gt}$.[16] To study the performance of the two-step estimator in estimating parameters of an assumed equilibrium, we start by computing the fixed point $\hat{\mathbf{P}}^*_{gt}$. Under

[16] McKelvey and Palfrey (1995) claim that uniqueness of the LQRE applies for "almost all games." Dagsvik (2020) establishes necessary and sufficient conditions for unique and multiple QRE and provides useful algorithms for concluding on the number of QRE. We applied the algorithms of Dagsvik (2020) to examine the uniqueness of the equilibrium in the eight games in duopoly and found that the equilibrium is unique for all eight games in duopoly.

Table 5
Confidence intervals from Monte Carlo simulation.

Model	Quadratic utility		Cobb–Douglas utility	
	95% C.I.[a]		95% C.I.[a]	
Preference parameters				
β_{Π}	1.05	1.05	0.67	0.71
β_{B}	0.45	0.70	0.29	0.33
$\beta_{\Pi\Pi}$	−0.05	−0.05	n.a.	n.a.
β_{BB}	−0.04	−0.02	n.a.	n.a.
$\beta_{\Pi B}$	−0.05	−0.04	n.a.	n.a.
$\underline{U}$	n.a.	n.a.	1.77	1.98
λ Parameters				
$Monopoly$	0.41	0.69	1.40	2.24
$Duopoly$	2.17	2.68	2.02	2.35
$Quadropoly$	3.76	4.55	2.31	2.73
Log-Likelihood	−6121.90		−6243.36	
# subjects:	136			
# Games:	24			

[a]95% confidence interval based on the Monte Carlo simulations described in Section 5.3.
Parameter normalizations: $\beta_{\Pi} + \beta_{\Pi\Pi} = 1$ in the quadratic specification, and $\beta_{\Pi} + \beta_{B} = 1$ in the Cobb–Douglas specification.

QRE, each player's belief of the opponent(s)'s choice probabilities is identical to the equilibrium probabilities. Hence, for a given occasion gt, the fixed point $\hat{\mathbf{P}}^{*}_{gt}$ is the solution to the following set of J non-linear equations:

$$\hat{P}_{1gt} = \frac{e^{\hat{\lambda}_t V_{1gt}(\hat{\mathbf{P}}^{*}_{gt})}}{\sum_{r\in C} e^{\hat{\lambda}_t V_{rgt}(\hat{\mathbf{P}}^{*}_{gt})}}\ ,\ \hat{P}_{2gt} = \frac{e^{\hat{\lambda}_t V_{2gt}(\hat{\mathbf{P}}^{*}_{gt})}}{\sum_{r\in C} e^{\hat{\lambda}_t V_{rgt}(\hat{\mathbf{P}}^{*}_{gt})}}\ ,\ \ldots, \hat{P}_{Jgt} = \frac{e^{\hat{\lambda}_t V_{Jgt}(\hat{\mathbf{P}}^{*}_{gt})}}{\sum_{r\in C} e^{\hat{\lambda}_t V_{rgt}(\hat{\mathbf{P}}^{*}_{gt})}}\ , \tag{13}$$

where, $\hat{P}_{jgt}$, the probabilities on the left-hand side of the equations are the J elements of $\hat{\mathbf{P}}^{*}_{gt}$.[17] After having computed the fixed point, we generate 5 000 synthetic data sets for both functional forms. We then simulate choices by individuals with preference and scale parameters given by $(\hat{\beta},\ \hat{\lambda}_t)$ under the assumption that they play the LQRE in (13). When generating the synthetic choice data, we draw type I extreme value distributed error terms, and let simulated players pick the treatment alternative, which maximizes:

$$\hat{\lambda}_t V_{jgt}(\hat{\mathbf{P}}^{*}_{gt}) + \varepsilon_{jgt}. \tag{14}$$

Using the 5 000 synthetic data sets where the data generating process is known, we estimate preference and scale parameters $(\hat{\beta}_1,\ \hat{\lambda}_{t1})\ldots(\hat{\beta}_{5000},\ \hat{\lambda}_{t5000})$ by applying the two-step procedure where we maximize (12). For each synthetic data set, we compute the observed relative frequencies of actions in a game and apply the program `gmnl` by Gu et al. (2013) to estimate 5 000 scaled logit models for both functional forms. We describe the distribution of $\hat{\lambda}_t$ in Fig. 1 and report the confidence intervals based on the Monte Carlo simulations in Table 5. For both quadratic and Cobb–Douglas specifications, the differences between the confidence intervals in Table 4 and those in Table 5 are negligible. The interpretation is that the reported standard errors of the two-step estimator are accurate.

The distributions in Fig. 1 are kernel density plots of estimated scale parameters in the three markets based on the synthetic data sets. We see that they do not overlap for the quadratic specification, whereas a substantial overlap is observed for the Cobb–Douglas specification. Hence, under the assumption that the quadratic model is the correct specification, it is highly unlikely that one would draw the wrong conclusion based on results from the two-step estimator, even with a small sample size. However, if the Cobb–Douglas model is the correct specification, one would need a larger sample size to provide statistical power to reject the null hypothesis that scale parameters in monopoly and duopoly are equal.

5.4. Model fit

As discussed by Wright and Leyton-Brown (2017), flexible models of human behavior are vulnerable for overfitting. Hence, the assessment of model fit should be based on out-of-sample predictions. By applying the following "jackknife-procedure," we show that estimation results from our QREC model can be used to predict behavior ex ante: We first estimate the model 24 times, excluding one game each time. Then we compute the predicted fixed point in the game that was excluded from the estimation. We denote these "jackknife fixed points" by $\hat{\mathbf{P}}^{*}_{-s}$, where $-s \in 1, 2, 3 \ldots 24$ denotes the excluded games.

In Fig. 2, we plot the observed relative frequencies for each strategy in each game, f_{jgt}, with the corresponding out-of-sample prediction, $\hat{\mathrm{P}}^{*}_{j,-s}$. Our QREC model captures the substantial differences in behavior across games and market settings quite well. The out-of-sample predictions are quite similar to the observed behavior. We do not reject the null hypothesis that the two distributions,

[17] See McKelvey and Palfrey, (1995), p. 11.

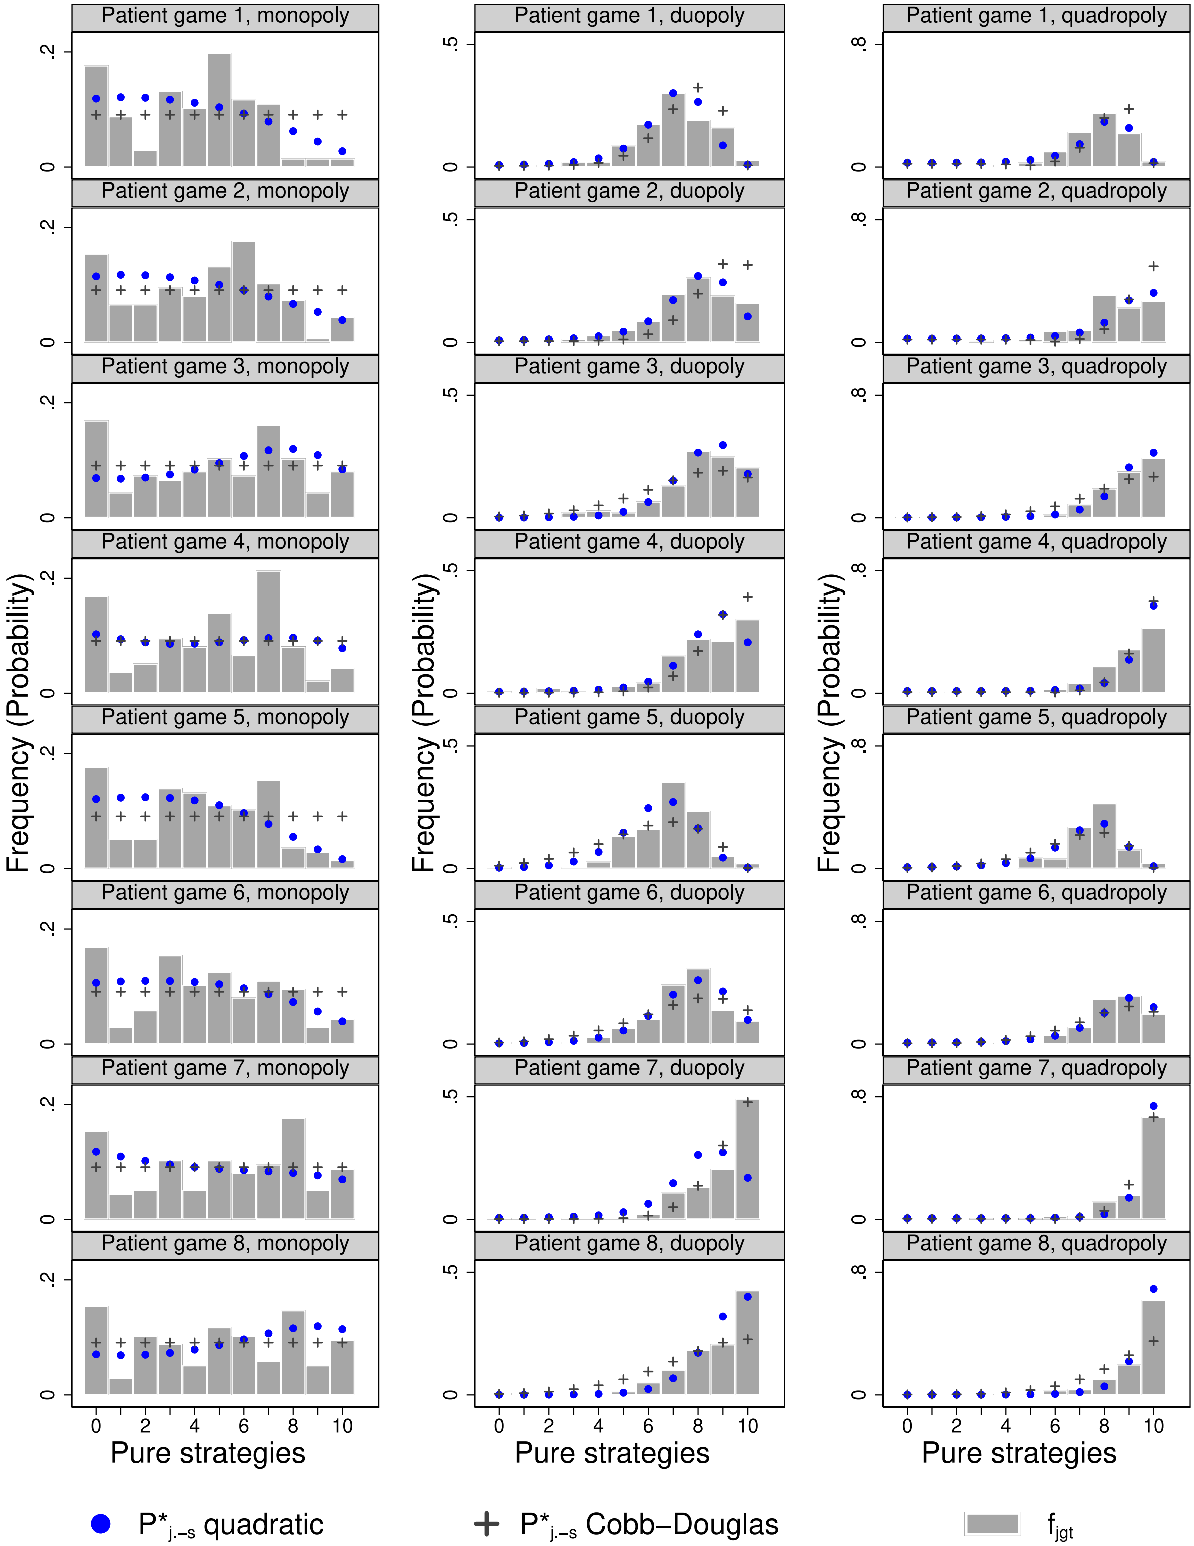

Fig. 2. Relative frequencies and fixed points predicted out of sample for all pure strategies in the experiment. Patient benefit in a game increases from left to right.

f_{jgt} and $\hat{P}^{*}_{j,-s}$, are the same by means of the Wilcoxon matched-pairs signed-ranks test (p = 0.1861 for quadratic, p = 0.1948 for Cobb–Douglas) and Fisher–Pitman permutation test for paired replicates (p > 0.99 for both quadratic and Cobb–Douglas forms). Judged by the mean squared error (MSE), the quadratic specification (MSE = 0.0021) provides a better fit than the Cobb–Douglas specification (MSE = 0.0033).

6. Discussion and conclusion

In this paper, we apply experimental data to study strategic medical choices in three market settings: monopoly, duopoly, and quadropoly. In our experiment, the individuals' choices result in more benefit for patients when there is more competition. Our results show that a substantial change in behavior can be explained by the fixed preference QREC model. The observed difference in behavior between monopoly, duopoly, and quadropoly market settings can be attributed to changes in the individuals' scale parameter while keeping preferences fixed. The scale parameter is a measure of randomness in behavior, or equivalently, a measure of determinism in behavior. We are unaware of any previous studies that allow for the possibility that changes in competition intensity affect scale. We did not have any prior expectation with regard to whether competition would reduce or increase the degree of determinism in behavior. On the one hand, decision making might be perceived as more complicated when having to account for the actions of one or more competitors, and raising complexity could bring about *more* randomness in treatment choices. On the other hand, competition might encourage effort and trigger the decision-makers' attention, leading to *less* randomness in treatment choices in market settings with more competition. We find that the scale parameter rises as the market becomes more competitive, implying a higher degree of determinism in behavior.

As highlighted in the examples of Louviere and Eagle (2006), ignoring scale or assuming scale to be fixed across contexts can lead to misleading conclusions on the stability of preferences. Our results provide useful nuances to the recent literature on whether markets erode social responsibility; see, e.g., Falk and Szech (2013), Bartling et al. (2015), Sutter et al. (2020) or Bartling et al. (2021). Recent contributions to this literature apply laboratory experiments to identify the effects of competition. Observed changes in behavior across experimental market settings are interpreted as evidence of unstable preferences. An example is Falk and Szech (2013) who, without estimating preference or scale parameters, conclude that market competition causes experimental subjects to value the life of a mouse less, based on the observation that competition in the laboratory has unfortunate consequences for the mice. If we had assumed fixed scale and context-dependent preferences, the results would have suggested that competition causes preferences to change, and that individuals become more altruistic in market settings with more competition. It might have been tempting to conclude that *"Competition raises moral values"* or, alternatively, that *"Competition crowds in pro-social motivation."*

The economic literature comprises many contributions that reject the assumptions of perfect rationality and purely selfish behavior. Rejection of fundamental assumptions has been constructive in that the process has led to the rise of behavioral economics (Thaler, 1988; Thaler and Ganser, 2015). Behavioral economics has not yet succeeded in proposing a new overarching behavioral economic theory, and Thaler (2018) predicts that such a theory will never come to existence. One may argue that a more optimistic view is warranted. The seminal contribution by McKelvey and Palfrey (1995) revealed that the toolkit of choice modelers can be applied to the study of strategic behavior. Yet, applications of choice models in studying strategic decision making in games with vector-payoff have not reached its potential. One may argue that empirical game theory literature can benefit from knowledge translation from advances in the vast choice modeling literature (Revelt and Train, 1998; McFadden, 2001; Hess and Daly, 2010; Jessie and Saari, 2016), where context-dependent scale parameters have been a central research topic for decades (Louviere et al., 1999, 2002; Louviere and Eagle, 2006; Louviere and Meyer, 2008; Wang et al., 2020). The QREC model, which is only a slight augmentation of a generalized multinomial logit model, provides a consistent economic model of strategic behavior that relies on less restrictive assumptions, thereby illustrating that choice modeling has a role in contributing to an overarching theory of economic behavior.

An important aspect of our experiment is the large number of cells in payoff matrices; 11^2 in duopoly and 11^4 in quadropoly. As a result, many potential choice combinations will remain unrealized in the data. Therefore, one may argue that the experimental design provides unfavorable conditions for maximum likelihood estimation of the parameters of a QREC model. However, results from the Monto Carlo simulations show that the two-step estimator produces accurate estimates with a moderate sample size, even under these unfavorable conditions. Furthermore, since the two-step estimator does not rely on the assumption that an equilibrium is reached, it is robust with regard to a hypothetical scenario where the data are *not* generated by a QRE.

Over the last decades, numerous laboratory experiments have generated a large base of data from, for example, public good games and ultimatum games. The fact that these games constitute multi-criteria games for altruistic decision-makers are typically ignored, and NE or other equilibrium concepts for single-criterion games are applied in the analyses. Re-examination of the existing data by fitting QREC models has substantial research potential. Behaviors that have been referred to as anomalies can be consistent with a QREC model with a plausible utility function. Hence, re-examination of the data is likely to contribute to new insights and provide explanations for phenomena that have been poorly understood. More methodological research on QREC models would benefit this proposed research. For example, while the two-step estimator is robust and consistent, whether it is more or less efficient than the full information maximum likelihood estimator described by Moffatt (2015) is a question for future research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank four anonymous reviewers for comments, questions, and suggestions which have contributed to substantial improvements of this paper. In our work with this paper, we have benefited much from inspiring scholarly conversations with John K. Dagsvik and Albert Ma. We thank Daniel Wiesen for joint collaboration on experimental design in the early stages of this project and for managing the experimental sessions in Cologne. We thank Emanuel Castillo for z-Tree programming and Anne Classen, Tobias Danzeisen, Lena Kuhne and Theresa Schwan for research assistance. We are grateful for comments, suggestions, and encouragement from Alexander W. Cappelen, Fredrik A. Dahl, Karen E. Hauge, Arne Risa Hole, Snorre Kverndokk, Eric Nævdal, Erik Ø. Sørensen, and Knut R. Wangen. We thank participants in seminars at The Ragnar Frisch Centre for Economic Research; FAIR, Norwegian School of Economics; School of Mathematical Sciences, University of Science and Technology, Hefei, China; and Economics and Management School, Wuhan University, China, for their questions and comments. We gratefully acknowledge financial support from the Research Council of Norway who provided the funds for designing the experiment, development of z-Tree program, conducting experiments at the University of Cologne, as well as writing and publishing the manuscript.

Funding

Research Council of Norway, IRECOHEX, No. 231776 Research Council of Norway, NORCHER, No. 296114.

Ethical review and approval of experimental procedure

Norwegian Social Science Data Services (reference No. 43709).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jocm.2021.100282.

References

Anderson, S.P., Goeree, J.K., Holt, C.A., 2001. Minimum-effort coordination games: Stochastic potential and logit equilibrium. Games Econom. Behav. 34 (2), 177–199.

Bajari, P., Hortacsu, A., 2005. Are structural estimates of auction models reasonable? Evidence from experimental data. J. Political Econ. 113 (4), 703–741.

Bartling, B., Fehr, E., Özdemir, Y., 2021. Does market interaction erode moral values? Rev. Econ. Stat. (forthcoming).

Bartling, B., Weber, R.A., Yao, L., 2015. Do markets erode social responsibility? Q. J. Econ. 130 (1), 219–266.

Battese, G.E., 1997. A note on the estimation of cobb-douglas production functions when some explanatory variables have zero values. J. Agric. Econ. 48, 250–252.

Blavatskyy, P.R., 2007. Stochastic expected utility theory. J. Risk Uncertain. 34 (3), 259–286.

Camerer, C.F., 2011. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press.

Dagsvik, J.K., 2008. Axiomatization of stochastic models for choice under uncertainty. Math. Social Sci. 55 (3), 341–370.

Dagsvik, J.K., 2015. Stochastic models for risky choices: A comparison of different axiomatizations. J. Math. Econom. 60, 81–88.

Dagsvik, J.K., 2016. What independent random utility representations are equivalent to the IIA assumption? Theory and Decision 80 (3), 495–499.

Dagsvik, J.K., 2018. Invariance axioms and functional form restrictions in structural models. Math. Social Sci. 91, 85–95.

Dagsvik, J.K., 2020. On the equilibria in models of social interaction and quantal response equilibrium. In: HERO Working Paper Series, vol. 2020:5, University of Oslo, Health Economics Research Network.

Erlander, S., 1998. Efficiency and the logit model. Ann. Oper. Res. 82, 203–218.

Falk, A., Szech, N., 2013. Morals and markets. Science 340, 707–711.

Fiebig, D.G., Keane, M.P., Louviere, J., Wasi, N., 2010. The generalized multinomial logit model: accounting for scale and coefficient heterogeneity. Mark. Sci. 29 (3), 393–421.

Fischbacher, U., 2007. Z-tree: Zurich toolbox for readymade economic experiments. Exp. Econ. 10, 171–178.

Gaynor, M., Town, R.J., 2011. Competition in health care markets. In: Handbook of Health Economics, vol. 2. Elsevier, pp. 499–637.

Gaynor, M., Vogt, W.B., 2000. Antitrust and competition in health care markets. In: Handbook of Health Economics, vol. 1. Elsevier, pp. 1405–1487.

Ge, G., Godager, G., 2021. Data from an incentivized laboratory experiment on strategic medical choices. Data in Brief.

Goeree, J.K., Holt, C.A., 2001. Ten little treasures of game theory and ten intuitive contradictions. Amer. Econ. Rev. 91 (5), 1402–1422.

Goeree, J.K., Holt, C.A., 2005. An experimental study of costly coordination. Games Econom. Behav. 51 (2), 349–364.

Goeree, J.K., Holt, C.A., Laury, S.K., 2002. Private costs and public benefits: unraveling the effects of altruism and noisy behavior. J. Publ. Econ. 83 (2), 255–276.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2005. Regular quantal response equilibrium. Exp. Econ. 8 (4), 347–367.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2010. Quantal response equilibria. In: Behavioural and Experimental Economics. Springer, pp. 234–242.

Gu, Y., Hole, A.R., Knox, S., 2013. Fitting the generalized multinomial logit model in Stata. Stata J. 13 (2), 382–397.

Gul, F., Pesendorfer, W., 2006. Random expected utility. Econometrica 74 (1), 121–146.

Haile, P.A., Hortaçsu, A., Kosenok, G., 2008. On the empirical content of quantal response equilibrium. Amer. Econ. Rev. 98 (1), 180–200.

Harsanyi, J.C., 1973. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. Internat. J. Game Theory 2 (1), 1–23.

Harsanyi, J.C., Selten, R., et al., 1988. A general theory of equilibrium selection in games. In: MIT Press Books, vol. 1. The MIT Press.

Hennig-Schmidt, H., Selten, R., Wiesen, D., 2011. How payment systems affect physicians' provision behavior – An experimental investigation. J. Health Econ. 30, 637–646.

Hess, S., Daly, A., 2010. Choice Modelling: The State-of-The-Art and the State-Of-Practice: Proceedings from the Inaugural International Choice Modelling Conference. Emerald Group Publishing Limited.

Hess, S., Rose, J.M., 2012. Can scale and coefficient heterogeneity be separated in random coefficients models? Transportation 39 (6), 1225–1239.

Hess, S., Train, K., 2017. Correlation and scale in mixed logit models. J. Choice Model. 23, 1–8.

Holte, J.H., Sivey, P., Abelsen, B., Olsen, J.A., 2016. Modelling nonlinearities and reference dependence in general practitioners' income preferences. Health Econ. 25 (8), 1020–1038.

Jessie, D.T., Saari, D.G., 2016. From the luce choice axiom to the quantal response equilibrium. J. Math. Psych. 75, 3–9.

Ke, S., 2018. Rational expectation of mistakes and a measure of error-proneness. Theor. Econ. 13 (2), 527–552.

Kirchler, M., Huber, J., Stefan, M., Sutter, M., 2015. Market design and moral behavior. Manage. Sci. 62 (9), 2615–2625.

Kolstad, J.R., 2011. How to make rural jobs more attractive to health workers. Findings from a discrete choice experiment in Tanzania. Health Econ. 20 (2), 196–211.

Loukopoulos, L.D., Dismukes, R., Barshi, I., 2009. The perils of multitasking. AeroSafety World 4 (8), 18–23.

Louviere, J.J., Eagle, T., 2006. Confound it! That pesky little scale constant messes up our convenient assumptions. In: Sawtooth Software Conference. Sawtooth Software Inc.

Louviere, J.J., Meyer, R.J., 2008. Formal choice models of informal choices. In: Review of Marketing Research. Emerald Group Publishing Limited, pp. 3–32.

Louviere, J.J., Meyer, R.J., Bunch, D.S., Carson, R., Dellaert, B., Hanemann, W.M., Hensher, D., Irwin, J., 1999. Combining sources of preference data for modeling complex decision processes. Mark. Lett. 10 (3), 205–217.

Louviere, J., Street, D., Carson, R., Ainslie, A., Deshazo, J., Cameron, T., Hensher, D., Kohn, R., Marley, T., 2002. Dissecting the random component of utility. Mark. Lett. 13 (3), 177–193.

Luce, R.D., 1959. Individual Choice Behavior a Theoretical Analysis. John Wiley, Oxford, England.

Ma, C.-t.A., Manove, M., 1993. Bargaining with deadlines and imperfect player control. Econometrica 1313–1339.

Mackowiak, P.A., 2020. George Washington's recurrent health problems and fatal infection reexamined. Clin. Infect. Dis. ciaa1324.

Matějka, F., McKay, A., 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. Amer. Econ. Rev. 105 (1), 272–298.

McCubbins, M.D., Turner, M., Weller, N., 2013. Testing the foundations of quantal response equilibrium. In: International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction. Springer, pp. 144–153.

McFadden, D., 1974. The measurement of urban travel demand. J. Publ. Econ. 3 (4), 303–328.

McFadden, D., 1981. Econometric models of probabilistic choice. In: Manski, C., McFadden, D. (Eds.), Structural Analysis of Discrete Data with Econometric Applications. MIT Press, pp. 198–272.

McFadden, D., 2001. Economic choices. Amer. Econ. Rev. 91 (3), 351–378.

McFadden, D., Machina, M.J., Baron, J., 1999. Rationality for economists? In: Fischhoff, B., Manski, C. (Eds.), Elicitation of Preferences. Springer, pp. 73–110.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econom. Behav. 10 (1), 6–38.

Moffatt, P.G., 2015. Experimetrics: Econometrics for Experimental Economics. Macmillan International Higher Education.

Nash, J.F., 1950. Equilibrium points in n-person games. Proc. Natl. Acad. Sci. 36 (1), 48–49.

Requate, T., Waichman, I., 2011. "A profit table or a profit calculator?" A note on the design of Cournot oligopoly experiments. Exp. Econ. 14 (1), 36–46.

Revelt, D., Train, K., 1998. Mixed logit with repeated choices: Households' choices of appliance efficiency level. Rev. Econ. Stat. 80 (4), 647–657.

Sarrias, M., Daziano, R., 2017. Multinomial logit models with continuous and discrete individual heterogeneity in R: The gmnl package. J. Stat. Softw., Articles 79 (2), 1–46.

Shapley, L.S., Rigby, F.D., 1959. Equilibrium points in games with vector payoffs. Nav. Res. Logist. Q. 6 (1), 57–61.

Sutter, M., Huber, J., Kirchler, M., Stefan, M., Walzl, M., 2020. Where to look for the morals in markets? Exp. Econ. 23 (1), 30–52.

Swait, J., Louviere, J., 1993. The role of the scale parameter in the estimation and comparison of multinomial logit models. J. Mark. Res. 30 (3), 305–314.

Swait, J., Marley, A.A., 2013. Probabilistic choice (models) as a result of balancing multiple goals. J. Math. Psych. 57 (1–2), 1–14.

Thaler, R.H., 1988. Anomalies: The ultimatum game. J. Econ. Perspect. 2 (4), 195–206.

Thaler, R.H., 2018. Behavioral economics: Past, present, and future. In: The Nora and Edward Ryerson Lecture at the University of Chicago. University of Chicago.

Thaler, R.H., Ganser, L., 2015. Misbehaving: The Making of Behavioral Economics. WW Norton New York.

Thurstone, L.L., 1927a. A law of comparative judgment. Psychol. Rev. 34 (4), 273.

Thurstone, L.L., 1927b. Psychophysical analysis. Am. J. Psychol. 38 (3), 368–389.

Train, K., 2009. Discrete Choice Methods with Simulation. Cambridge University Press.

Tversky, A., 1972. Choice by elimination. J. Math. Psych. 9 (4), 341–367.

van der Pol, M., Currie, G., Kromm, S., Ryan, M., 2014. Specification of the utility function in discrete choice experiments. Value Health 17 (2), 297–301.

Van Der Pol, M., Shiell, A., Au, F., Jonhston, D., Tough, S., 2010. Eliciting individual preferences for health care: a case study of perinatal care. Health Expect. 13 (1), 4–12.

Voorneveld, M., Grahn, S., Dufwenberg, M., 2000. Ideal equilibria in noncooperative multicriteria games. Math. Methods Oper. Res. 52 (1), 65–77.

Voorneveld, M., Vermeulen, D., Borm, P., 1999. Axiomatizations of Pareto equilibria in multicriteria games. Games Econom. Behav. 28 (1), 146–154.

Wallin, A., Swait, J., Marley, A., 2018. Not just noise: A goal pursuit interpretation of stochastic choice. Decision 5 (4), 253.

Wang, J., Iversen, T., Hennig-Schmidt, H., Godager, G., 2020. Are patient-regarding preferences stable? Evidence from a laboratory experiment with physicians and medical students from different countries. Eur. Econ. Rev. 125.

Wright, J.R., Leyton-Brown, K., 2017. Predicting human behavior in unrepeated, simultaneous-move games. Games Econom. Behav. 106, 16–37.

Zeleny, M., 1975. Games with multiple payoffs. Internat. J. Game Theory 4 (4), 179–191.