

QRE vs Nash equilibrium: Evaluating the Jury Paradox in Voting Games

April Cates

A dissertation submitted to Auckland University of Technology in
partial fulfilment of the requirements for the degree of Master of
Business (MBus)

March 6, 2025

Faculty of Business, Economics, and Law

Supervisor: Dr. Matthew Ryan

Abstract

In traditional game theory, Nash Equilibrium (NE) has been the dominant framework for predicting and explaining behaviour in strategic settings, including voting games. However, NE assumes fully rational decision-making in the sense that it assumes individuals make choices which are free of error. Yet, this often fails to capture real-world behaviour. This limitation is particularly evident in the jury paradox, where rational jurors, following their strategic incentives, may systematically convict innocent defendants or fail to convict the guilty. Empirical evidence suggests that, in reality, jurors do not strictly adhere to Nash equilibrium predictions, particularly when uncertainty and asymmetric payoffs come into play.

This dissertation examines the prevalence and persistence of the jury paradox under both a Nash Equilibrium and Quantal Response Equilibrium (QRE) framework. QRE introduces bounded rationality by allowing for probabilistic decision-making, which better reflects real-world uncertainty and its impact on voting behaviour. Using computational simulations, we analyse how equilibrium outcomes shift as decision noise, payoff asymmetry, and uncertainty increase. Our findings suggest that QRE not only captures deviations from Nash equilibrium behaviour but also provides a more accurate representation of jury voting dynamics, particularly in cases where Nash equilibrium predictions exhibit a jury paradox.

The results highlight the importance of incorporating bounded rationality into equilibrium models to accurately reflect real-world decision-making. This research deepens our understanding of how the jury paradox emerges and persists across different equilibrium models, offering insights relevant to both theoretical and applied game theory.

Contents

1	Introduction	7
2	The Model	9
3	Symmetric strategies and equilibria	11
4	Equilibria and best responses under a unanimous rule	12
5	A visualisation of equilibria and best responses under a unanimous rule	14
5.1	Further comparative statics	16
6	Error Probabilities	20
6.1	The Jury Paradox	21
7	Quantal Response Equilibrium	24
7.1	Quantal Response Equilibrium in this model	25
8	Logit QRE	28
9	QRE and The Jury Paradox	28
10	Results and Discussion	29
11	Significance and Limitations	32
12	Conclusion	33
13	Appendix A: Jury Paradox Comparison When $c = \frac{1}{9}$	36
14	Appendix B: Jury Paradox Comparison When $c = \frac{2}{9}$	37

List of Figures

1	NE indifference curves under unanimity	16
2	Indifference and contour curves under unanimity	21
3	QRE feasibility regions $n = 3$	27
4	QRE feasibility regions $n = 6$	27
5	Type-1 error curve for different λ and NE	32

List of Tables

1	Type-1 error rates for different hurdles when $n = 6$.	22
2	Type-1 error rates for $n = 6$ and $c = \frac{2}{9}$.	23
3	Type-1 error rates for $n = 6$ and $h = \{5, 6\}$.	24
4	NE (c, r) pairs for $n = 6$.	30
5	QRE (c, r) pairs for $\lambda = 10$.	31
6	QRE (c, r) pairs for $\lambda = 25$.	31
7	Table for $n = 3$.	36
8	Table for $n = 4$.	36
9	Table for $n = 5$.	36
10	Table for $n = 6$.	36
11	Table for $n = 7$.	36
12	Table for $n = 8$.	36
13	Table for $n = 9$.	36
14	Table for $n = 10$.	36
15	Table for $n = 11$.	37
16	Table for $n = 12$.	37
17	$n = 4, c = \frac{2}{9}$.	37
18	$n = 5, c = \frac{2}{9}$.	37
19	$n = 6, c = \frac{2}{9}$.	37
20	$n = 7, c = \frac{2}{9}$.	37
21	$n = 8, c = \frac{2}{9}$.	37
22	$n = 9, c = \frac{2}{9}$.	37
23	$n = 10, c = \frac{2}{9}$.	38
24	$n = 11, c = \frac{2}{9}$.	38
25	$n = 12, c = \frac{2}{9}$.	38

Attestation

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor used artificial intelligence tools or generative artificial intelligence tools (unless it is clearly stated, and referenced, along with the purpose of use), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Signature:

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor, Dr. Matthew Ryan, for his invaluable insight, support, and patience throughout the process of researching and writing this dissertation. I would also like to thank my mum and dad, for whom I am eternally grateful. Your belief in me is the only reason this dissertation was possible. I want you to know that your unconditional support and love means more to me than you will ever know. I also want to thank my brother, whose ability to make light out of even the most stressful situations sustained and lifted me when I needed it most. I also want to acknowledge the contributions of my late grandparents, who always encouraged me to pursue higher education. Last but certainly not least, I want to thank my cat for keeping me sane. I hope I can make you all proud. Thank you.

1 Introduction

The assumption that individuals make optimal decisions was once an unchallenged notion among traditional game theorists. For instance, within the field of standard or traditional game theory, the assumption that human beings are perfect optimizers or are perfectly rational underpins a great deal of the existing theories and literature (Goeree, Holt, and Palfrey 2020). However, a contrary assumption underpins the well-known Condorcet jury theorem. In this theorem, it is assumed that an individual will behave in an identical manner regardless of whether or not they are voting as part of a group or on their own to determine an outcome (Condorcet 1785).

The Condorcet jury theorem shows that in a binary choice scenario, if each juror has a better-than-even chance of making the "correct" choice, given the independent information they've received, then the probability that a verdict reached by majority rule will obtain the correct result approaches one as the size of the jury goes towards infinity. It is important to note that the proofs which underpin the Condorcet jury theorem assume that each juror votes informatively, or in other words, votes according to the private innocent or guilty signal they receive. However, later research, such as Austen-Smith and Banks (1996), points out that such behaviour is often sub-optimal and is inconsistent with the general intuition surrounding collective decision-making and Nash equilibria. Instead, it is suggested that jurors vote in a strategic manner, where they consider not only their individual private signals but also the behaviour of their peers. The experimental findings, such as those by Bouton, Llorente-Saguer, and Malherbe (2017) prove that strategic voting is indeed relatively good at predicting individual behaviour.

Yet in contrast to Condorcet's assumption that decision accuracy improves with jury size, Feddersen and Pesendorfer (1998) find that this is not always the case and is greatly dependant on the voting rule. In particular, they find that under unanimity, innocent defendants are often convicted at a higher rate than under any other voting rule, and this rate does not approach zero as the jury size is increased. This finding goes against the long-held belief that the requirement of a unanimous verdict offers greater protection to the innocent than any other voting rule (Coughlan 2000). Experimental evidence, such as that gathered by Guarnaschelli, McKelvey, and Palfrey (2000), further highlights the inconsistency of Condorcet's assumption through the large number of incorrect decisions observed.

Table 7 from Guarnaschelli, McKelvey, and Palfrey (2000) offers a clear example of why quantal response equilibrium (QRE) may provide us with a more suitable way to analyse data and assess the jury paradox. As previously mentioned, the table exhibits a Nash equilibrium jury paradox for $n = 6$. In the following QRE column, we are shown that the QRE is actually able to eliminate this paradox as it reports a reduction in the proportion of incorrect convictions from 0.187

to 0.069 as the hurdle required for conviction moves from four to six. However, Table 7 cannot say with certainty whether QRE is better suited to handle the paradox as Guarnaschelli, McKelvey, and Palfrey (2000) vary a parameter in the QRE model when varying the conviction hurdle in order to ensure that the model provides a good fit for their data. We shall delve into the specifics of this as the paper progresses.

Another instance of a jury paradox is exhibited in Table 1 from the paper by Feddersen and Pesendorfer (1998). In this table, we are shown that in a twelve-person jury, as the hurdle required to convict increases from seven to twelve, so does the probability of convicting an innocent defendant. It is interesting to note here that Feddersen and Pesendorfer (1998) discovered that for an infinitely large jury, unanimity always exhibits a positive Type-1 error, whereas any hurdle below unanimity converges to zero Type-1 error as the jury becomes larger. However, the case of finite-sized juries is not as straightforward. For instance, juries with a finite n may or may not exhibit a jury paradox depending on the other parameters.

The above finding implies that we may need to search for other game-theoretic tools to better analyse the data. Quantal response equilibrium (QRE) developed by McKelvey and Palfrey (1995) provides us with a suitable alternative due to its flexibility and ability to better explain real-world behaviour.

The previously mentioned finding from Feddersen and Pesendorfer (1998), in which, under unanimity, the probability of convicting an innocent defendant does not decrease as the hurdle for conviction increases, is an example of a "jury paradox." The probability of convicting an innocent defendant will be described in this dissertation as the Type-1 error probability. The Type-1 error probability is an important metric to gather and analyse as it directly influences the "jury paradox." As already discussed, the jury paradox refers to any instance in which the Type-1 error probability does not decrease as the hurdle required for conviction increases. Another way of saying this is that a jury paradox will occur if an increase in the conviction hurdle also results in an increase in the probability of convicting an innocent defendant. To give a concrete example, a Nash equilibrium jury paradox can be observed in Table 7 from the paper by Guarnaschelli, McKelvey, and Palfrey (2000) for $n = 6$ but not $n = 3$. To explain why this is so, let's first consider the case where $n = 6$. Under majority, where at least four jurors must vote to convict, the probability that an innocent defendant is convicted is equal to 0.71. However, under unanimity, this probability rises to 0.186. Therefore, we can conclude that a jury paradox is evident in this case as the Type-1 error has increased rather than decreased with the hurdle. Now, we can analyse the scenario where $n = 3$ to uncover why there is no jury paradox evident for a jury of this size. Under majority, where at least two jurors must vote to convict, the probability that an innocent defendant is convicted is equal to 0.216. As we move to unanimity, this probability falls to 0.140. Given that a shift from a simple majority to a unanimous voting rule

has reduced the Type-1 error probability, we can conclude that a jury paradox does not exist in this scenario.

This paper aims to explore the prevalence of the jury paradox under Nash equilibrium versus QRE, and the factors that contribute to the existence and pervasiveness of the paradox. The majority of the existing literature only compares Nash equilibria to QRE for symmetric payoffs in which both errors have identical payoffs. Therefore, this paper aims also to contribute by assessing the qualities and pervasiveness of the paradox under asymmetric payoffs.

2 The Model

It is assumed that there are two states of the world, one where the defendant is innocent and the other where the defendant is guilty. The true state is denoted by $s \in S = \{[I], [G]\}$ where $s = [I]$ represents the state where the defendant is innocent and $s = [G]$ represents the state where the defendant is guilty. The jury is composed of n jurors who will decide to vote either to acquit or convict the defendant. Each juror shares a common prior regarding the state of the world, which is represented by $Pr[s = [I]] = p \in (0, 1)$. Prior to voting, each juror will receive a private signal $t \in \{I, G\}$ which is independently and identically distributed, conditional on the true state of the world. The probability that jurors receive either a guilty or innocent signal given the true state of the world is denoted as $Pr[[t] = s|s] = r \in (0.5, 1)$. Without communicating with their peers, jurors will then finally cast their vote $v \in \{I, G\}$ to either vote innocent or guilty. The outcome of the vote $o \in \{A, C\}$ will result in either an acquittal, A , or a conviction, C . Under any voting rule, a defendant will be convicted if and only if the hurdle required for conviction is met or exceeded. This hurdle shall be denoted as h such that a defendant will only be convicted if and only if at least h jurors vote guilty.

All jurors share a common Bernoulli utility payoff, which is dependent on the outcome of the vote and the state of the world, such that they will receive the following payoffs, where $c \in (0, 1]$:

$$v(A, [I]) = v(C, [G]) = 0$$

$$v(A, [G]) = -c$$

$$v(C, [I]) = -1$$

These payoffs differ from those presented by Feddersen and Pesendorfer (1998) in which jurors receive the following payoffs,

$$u(A, [I]) = u(C, [G]) = 0$$

$$u(C, [I]) = -q$$

$$u(A, [G]) = -(1 - q)$$

where $q \in (0, 1)$. The parameter q is used to represent a threshold of reasonable doubt. This means that in order for jurors to prefer conviction over acquittal, the belief in a defendant's guilt must exceed this parameter q .

It can be shown that for each value of c , there is a value for q where the Feddersen and Pesendorfer (1998) payoff functions are a positive linear transformation of those used in the described model. This means that the Feddersen and Pesendorfer (1998) payoff functions represent the same preferences as those contained within this model. The linear mapping calculations are detailed below, using the formula $T(x) = ax + b$. The following payoffs need to be mapped:

$$v(A, [I]) = v(C, [G]) = 0 \text{ to } u(A, [I]) = u(C, [G]) = 0$$

$$v(A, [G]) = -c \text{ to } u(A, [G]) = -(1 - q)$$

$$v(C, [I]) = -1 \text{ to } u(C, [I]) = -q$$

Given that the first set of payoffs are both equal to zero, the general equation $T(x) = ax + b$ can be simplified further.

$$T(0) = a(0) + b$$

$$T(0) = b = 0$$

So the general formula can then be re-written as:

$$T(x) = ax$$

Now this simplified formula can be used to map the remaining utility payoffs. First by setting x equal to $-c$

$$T(-c) = a(-c) = -(1 - q)$$

$$a = \frac{1 - q}{c}$$

Then lastly, setting x equal to -1 to obtain

$$T(-1) = a(-1) = -q$$

$$a = q$$

It is important to note here that $a > 0$. Given that there are two expressions for a we can set them equal to each other to solve for the value of q :

$$\frac{1 - q}{c} = q$$

$$1 - q = qc$$

$$1 = q(c + 1)$$

$$q = \frac{1}{c + 1}$$

Which results in the final linear mapping equation $T(x) = \frac{1}{c+1}(x)$.

The parameter c plays an important role in determining a juror's attitude towards risk. To better understand how c affects risk attitude, it is useful to provide some experimental context. For instance, imagine there is an experiment in which every player (juror) will earn a certain sum, $\$H$ if the jury arrives at a correct decision. If instead, the jury makes a Type-two error, where they acquit a guilty defendant, then they instead receive a different sum of $\$M$. Lastly, if the jury instead makes a Type-one error, where they convict an innocent defendant, then they shall receive a different sum again denoted as $\$L$. It is important to note that it is assumed that $\$H > \$M > \$L$. We can then normalise the utility so that $u(H) = 0, u(M) = -c$ and $u(L) = -1$. It is unreasonable to assume that all jurors share the same attitudes towards risk. Therefore, the value of c , and hence the expected utility of any lottery over $\{H, M, L\}$ will vary among jurors. To visualise how c impacts risk attitude, suppose jurors are offered two choices they may decide between. They can either choose to receive $\$M$ right away or partake in a lottery where they can win $\$H$ with probability p or $\$L$ with probability $1 - p$. We can then calculate for what value of p makes jurors indifferent between the lottery and the guaranteed value of $\$M$ to get the following: $c = 1 - p$. More risk-averse jurors require a greater value of p in order to be indifferent between the lottery and the guaranteed sum of $\$M$. A higher p also leads to a lower c value, which tells us that the lower a juror's c value, the more averse they are to risk, and hence, the more likely they are to choose the guaranteed sum over the lottery.

3 Symmetric strategies and equilibria

The behaviour of a single juror, i , can be represented as a strategy $\sigma^i = (\sigma_I^i, \sigma_G^i)$. This strategy encapsulates the probability of casting a guilty vote depending on whether juror i observes a guilty or an innocent signal. When all jurors use a common strategy, this is referred to as a symmetric profile. In this case, individual juror superscripts can be removed from the notation to represent a

general profile for all jurors $\sigma = (\sigma_I, \sigma_G)$. A symmetric equilibrium can then be identified as a Bayesian Nash equilibrium formed from a symmetric profile. Attention is restricted to symmetric profiles henceforth.

4 Equilibria and best responses under a unanimous rule

Suppose a given juror believes that every other juror votes informatively (i.e., the juror believes that their peers vote to convict if they receive a guilty signal and vote to acquit if they receive an innocent signal). Under the unanimity rule, we know that in order for this juror to be pivotal, all other jurors must have voted to convict the defendant, and hence, every other juror received a guilty signal. If a juror is pivotal, it means that their vote alone determines the final outcome. With this in mind, suppose we make the assumption that a juror who knew that every signal received by their peers was a guilty signal favours convicting the defendant even when their own signal is innocent. This assumption will hold if n is sufficiently large, with “sufficiency” depending on the value of r – higher r values reduce the required n . Given that the juror who receives an innocent signal prefers to vote guilty, conditional on pivotality and their own signal, informative voting cannot be an equilibrium. Therefore, given that under a unanimity rule, informative voting is not an equilibrium (at least for a large n) voting must occur in mixed strategies as discussed in the following section.

The parameter π_I represents the probability that the defendant is innocent, given a juror receiving an innocent signal. Conversely, π_G represents the probability that the defendant is innocent, given a juror receiving a guilty signal. Both π_I and π_G are calculated using Bayes’ rule such that:

$$\pi_I = \frac{rp}{rp + (1-r)(1-p)}$$

and

$$\pi_G = \frac{(1-r)p}{(1-r)p + r(1-p)}$$

The expected utility for each vote in relation to each signal is calculated below, where $\hat{\sigma}$ is used to denote the common strategy used by all other voters. Regardless of whether a juror receives a guilty or innocent signal, all jurors vote to maximise their expected utility, conditional on their vote being pivotal. The probability that a juror is pivotal conditional on each state is expressed using the notation z and y , where z represents the probability that a juror is pivotal when $s = [G]$, and y represents the probability that a juror is pivotal when $s = [I]$ such that:

$$z = C(r\sigma_G + (1-r)\sigma_I)^{h-1}(r(1-\sigma_G) + (1-r)(1-\sigma_I))^{n-h}$$

and

$$y = C(r\sigma_I + (1-r)\sigma_G)^{h-1}(r(1-\sigma_I) + (1-r)(1-\sigma_G))^{n-h}$$

where h is the hurdle required for conviction and C represents the number of ways of assigning $h-1$ of $n-1$ jurors to vote guilty.

The expected utility that results from voting innocent is equal to the probability that the defendant is actually guilty multiplied by the payoff associated with acquitting a guilty defendant ($-c$):

$$Eu(t, v = I) = (1 - \pi_t)(-c)$$

The expected utility associated with voting guilty is slightly more involved than the previous formula. This expected utility is the sum of two components. The first is the probability that a juror is not pivotal when the true state is guilty ($1-z$), multiplied by the probability of guilt ($1-\pi_t$), and all this is multiplied by the payoff associated with acquitting a guilty defendant ($-c$). The second multiplies the probability that a juror is pivotal when the true state is innocent (y), by the probability that the defendant is innocent (π_t), and all this is multiplied by the payoff associated with convicting an innocent defendant (-1).

$$Eu(t, v = G) = (1-z)(1-\pi_t)(-c) + (y)(\pi_t)(-1)$$

Now that we have the above formulas, we can calculate the indifference condition for both signals as follows:

$$\frac{1-z}{y} \frac{1-\pi_t}{\pi_t} c - 1 = 0$$

which can be rearranged and simplified to get:

$$\frac{1-z}{y} = \frac{\pi_t}{c(1-\pi_t)}$$

This value of indifference will be a function of the parameters within the model. For example, we can isolate π_I to solve for the value of indifference associated with the ratio of $\frac{\hat{\sigma}_I}{\hat{\sigma}_G}$:

$$Eu(t = I, v = G) - Eu(t = I, v = I) = 0$$

$$\left(\frac{((1-r)\hat{\sigma}_I + r\hat{\sigma}_G)}{r\hat{\sigma}_I + (1-r)\hat{\sigma}_G} \right)^{n-1} (c)(1-\pi_I) = \pi_I$$

Letting $x = \frac{\hat{\sigma}_I}{\hat{\sigma}_G}$ and hence $\hat{\sigma}_I = \hat{\sigma}_G x$, we can then substitute x into the previous equation for π_I to get:

$$\left(\frac{(1-r)x+r}{rx+(1-r)} \right)^{n-1} (c)(1-\pi_I) = \pi_I$$

Now that we have this expression for π_I we can use it to solve for x where:

$$x = \frac{\left(\frac{\pi_I}{c(1-\pi_I)} \right)^{\frac{1}{n-1}} (1-r) - r}{1-r - \left(\frac{\pi_I}{c(1-\pi_I)} \right)^{\frac{1}{n-1}} r}$$

This expression for x is based on the symmetric profile being responded to. This expression tells us that the best response for a juror who observes an innocent signal depends only on the ratio of the components of the other jurors' common strategy.

5 A visualisation of equilibria and best responses under a unanimous rule

Figure 1 is a helpful tool that can be used to identify and explain the mechanics of all possible Nash equilibria. Both the following graph and discussion refer to the case where $p = 0.5$, $n = 3$, $r = 0.7$, and $c = 1$. On the axes of the graph are the components of the common rival strategy to which the juror is responding. On the x-axis is the probability a single juror votes to convict upon receiving an innocent signal, and the y-axis represents the probability that a single juror votes to convict upon receiving a guilty signal. Both the red and the blue lines are indifference curves in the sense that they represent the points that make jurors indifferent between voting to convict or acquit upon receiving either a guilty or innocent signal, assuming that all other jurors vote in accordance with the indicated strategy. The blue line represents the indifference curve for when $t = I$ and the red line represents the indifference curve for when $t = G$. It is first helpful to familiarise ourselves with Figure 1, and how it can be used to identify all possible equilibria.

In order for any point in Figure 1 to be an equilibrium, it must be a best response to itself, meaning that if every other juror is using that strategy, it is also in the juror's best interest to use this same strategy. The first and perhaps most intuitive equilibrium is located at the point $(0, 0)$ in Figure 1. This point is often referred to as the trivial equilibrium. The trivial equilibrium arises when every other juror has voted to acquit, so the final juror can vote in whichever manner they please, as any response is a best response. At any point above

both indifference curves it is a best response to vote guilty as the expected utility exceeds that associated with voting to acquit. The opposite is then true for any point below both curves. For instance, consider a juror who is responding to a situation in which every other juror is utilising a strategy located above both indifference curves. In this case, a best response would be to vote guilty no matter what, which would lead you to the point $(1, 1)$ in Figure 1. However, this can not be the point that the juror is responding to as it is not above both curves, so we don't have an equilibrium here. Next, consider the points below both indifference curves. If a juror is responding to a situation in which every other juror is utilising a strategy located below both indifference curves, then a best response is to vote innocent no matter what. This will lead them to the point $(0, 0)$, which we know is the trivial equilibrium and does not lie below both curves. Finally, if a juror is now responding to a point in between both indifference curves, a best response would now be to vote guilty following a guilty signal and to vote innocent following an innocent signal. This best response is located at the point $(0, 1)$ in Figure 1. Again, however, this point can only be an equilibrium if this is the point that the juror is responding to. As $(0, 1)$ is not located in between the two indifference curves, we can deduce that this point is also not an equilibrium. Given that the non-trivial equilibrium is not located below, between, or above the indifference curves, we must now turn our attention to the curves themselves. At any point along the lines, jurors are indifferent between voting to convict following their signal. However, the only point along the blue line that is a best response to itself, and hence is an equilibrium, is located at the point at the top of the line $(0.3, 1)$. This point is one of the two equilibria that form the equilibrium set for Figure 1. The second is, of course, the trivial equilibrium located at the point $(0, 0)$. Turning now to the points along the red curve. The only point along the red curve that is a best response to itself, and hence is an equilibrium, is the trivial equilibrium located at the point $(0, 0)$. The particular structure of these equilibria is, of course, subject to change depending on the slopes of the indifference curves. For example, as the slope of the blue curve reduces, the equilibrium value of σ_I will increase as the top of the blue curve moves clockwise. While these other equilibrium structures are important to note, the ones previously discussed will be where we keep our attention. The reason for this is that this structure is used in the existing literature and hence allows for greater comparison.

To understand the mechanics and intuition behind the non-trivial equilibrium, suppose that every juror votes informatively. Now, suppose that a single juror has found themselves in a pivotal position after observing an innocent signal. When a juror becomes pivotal in an infinitely large jury, the sheer volume of guilty votes can cause the pivotal juror to doubt the reliability of their signal and consider it a possible outlier or error. This provides the pivotal juror with an incentive to deviate. Therefore, informative voting cannot be an equilibrium.

Now consider a scenario where jurors utilize mixed strategies such that they randomize following an innocent signal. In the above case of informative voting,

we showed that guilty votes offer compelling evidence that a juror did indeed receive a guilty signal. However, this is no longer the case when jurors vote using mixed strategies, as a portion of the guilty votes may have arisen from randomization rather than genuine informative voting. This possibility of randomization provides the pivotal juror with less of an incentive to deviate, as guilty votes are now less reliable indicators of guilt. Therefore, if σ_I is fixed at a value that makes jurors indifferent following an innocent signal, the equilibrium is held together as needed.

Indifference Curves for σ_I and σ_G under Unanimity

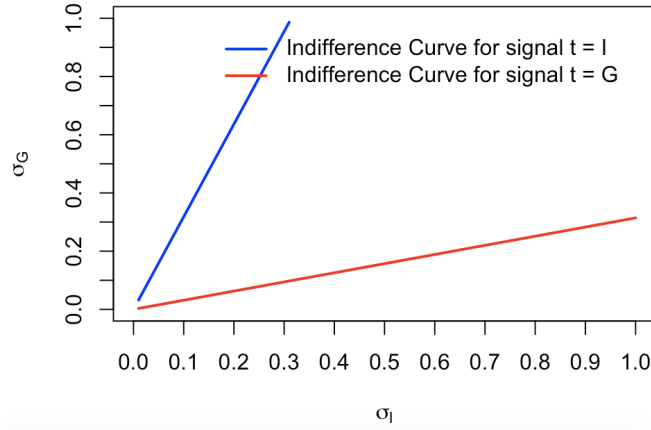


Figure 1: NE indifference curves under unanimity

5.1 Further comparative statics

The equilibria and slopes of the indifference curves pictured in Figure 1 are all influenced by the parameters n, p, r, c . By exploring such comparative statics, we can isolate the impact of each parameter on the indifference curves and equilibria. To isolate the impact of each individual parameter, we can take our indifference equation and solve it for σ_G as a function of σ_I to get the equation for each indifference curve. Given the indifference equation:

$$\left(\frac{(1-r)\sigma_I + r\sigma_G}{r\sigma_I + (1-r)\sigma_G} \right)^{n-1} = \frac{\pi_t}{c(1-\pi_t)}$$

we can solve and rearrange to get:

$$\sigma_G = \frac{l_t r - (1-r)}{r - l_t(1-r)} \sigma_I$$

where l_t is equal to $\frac{\pi_t}{c(1-\pi_t)^{n-1}}$ and is dependant on the parameters $n, c, r, \text{ and } p$. In order to examine the effects of each parameter on the slope of the indifference curves and the possible set of equilibria, let's refer back to the scenario in Figure 1 where $p = 0.5, r = 0.7, c = 1, n = 3$. The slopes of each indifference curve can range between both positive and negative values depending on the values of r and h . While the slopes of the curves can be either positive or negative, the red curve will always be positioned clockwise relative to the blue curve. This is because the region on the graph that corresponds to voting for guilty, given an innocent signal, is a subset of the region corresponding to voting guilty given a guilty signal. This can be shown by rearranging the payoff functions to prove that a juror will vote guilty when their expected utility from doing so exceeds that from voting innocent. For π_I :

$$(1 - \pi_I)c(1 - z) - y\pi_I > 0$$

For π_G :

$$(1 - \pi_G)c(1 - z) - y\pi_G > 0$$

The term $(1 - \pi_t)c(1 - z)$ is smaller for $t = I$ than for $t = G$ and the term $-y\pi_t$ is more negative for $t = I$ than for $t = G$. These findings indicate that of the above two inequalities, the one for $t = I$ is stricter and harder to satisfy than that for $t = G$, meaning that the region for voting guilty when $t = I$ is indeed a subset of voting for guilty when $t = G$. π_G will always be less than π_I for any $r > 0.5$ and for any p that is greater than 0 but less than 1. This means that the slope of the red curve should always be less than that of the blue curve, as long as the slope of the blue curve is positive. As the slopes of the curves vary, so do the possible set of equilibria as the parameter values change.

Let us now examine the comparative statics of each parameter when both curves have a positive slope. First, however, it is necessary to identify what parameter restrictions ensure that the slope of the indifference curve for $t = I$ is positive and greater than one. This is an important first step to take as these conditions enable a unique non-trivial equilibrium to occur. The slope of the indifference curve is denoted as follows:

$$Slope_I = \frac{l_I r - (1 - r)}{r - l_I(1 - r)}$$

Firstly, we must establish the conditions under which the slope is greater than zero. The slope of the indifference curve will be greater than zero under two circumstances. First, if both the denominator and numerator are greater than zero, and second, if both the denominator and numerator are less than zero. If either of those scenarios occurs, then we can conclude that the slope will be greater than zero. Numerically, we can denote this condition as $\frac{1-r}{r} < l_I < \frac{r}{1-r}$. Given that $r > 0.5$ we can rule out the possibility that both the numerator and denominator are negative so we must have that they are both positive and therefore that $l_I r - (1 - r) > 0$ and $r - l_I(1 - r) > 0$. Both of these inequalities can

be arranged to yield $l_I > \frac{1-r}{r}$ and $l_I < \frac{r}{1-r}$. By combining these inequalities, we can prove that $\frac{1-r}{r} < l_I < \frac{r}{1-r}$.

Now we must establish the conditions under which the slope is greater than one. The slope of the indifference curve will be greater than one as long as the numerator of the slope equation is larger than the denominator and both are greater than zero. We can calculate that the slope of the indifference curve for $t = I$ will be greater than one as long as $l_I > 1$ and $\frac{1-r}{r} < l_I < \frac{r}{1-r}$. The proof of this statement is as follows. We can write that we need the slope to be greater than one as $\frac{l_I r - (1-r)}{r - l_I(1-r)} > 1$. Given that $r - l_I(1-r) > 0$ we can multiply each side of the inequality to by this expression to get $l_I r - (1-r) > r - l_I(1-r)$ which, by rearranging and collecting like terms, can be re-expressed as $l_I - (1-r) > r$. Finally, by adding $(1-r)$ to each side, the expression simplifies further and becomes $l_I > r + (1-r)$. Finally, given that $r + (1-r) = 1$ we can conclude that the slope of the indifference curve will be greater than one as long as $\frac{1-r}{r} < l_I < \frac{r}{1-r}$. With these conditions in place we can begin to look into the comparative statics of each individual parameter, beginning with l_I . As l_I increases, so does the slope of the indifference curve as both the numerator and denominator are positive.

To examine the effect of altering the value of n we must instead look at the expression for l_I directly. As n is only contained within the exponent of the expression for l_I , we can calculate that as n increases, l_I will decrease since $l_I > 1$. Therefore, as n increases and l_I decreases, the slope of the indifference curve for $t = I$ will decrease. This reduction in slope causes the equilibrium value of σ_I to increase, which in turn increases the probability that jurors vote to convict following an innocent signal. A similar calculation can be done in order to determine the impact of c . As c is in the denominator of the fraction representing l_I we can see that as it increases, the value of l_I will decrease. We can therefore conclude that as c increases and l_I decreases, the slope of the indifference curve for $t = I$ will ultimately decrease.

Now, exploring the impact of p , we must look directly at the expression for π_I as p is only contained within this expression. By differentiating π_I with respect to p we get the following:

$$\frac{\partial \pi_I}{\partial p} = \frac{r(1-r)}{(rp + (1-r)(1-p))^2}$$

The denominator of the above derivative must always be positive given that it is squared, but the sign of the numerator depends on the value of p . Given that $0.5 < r < 1$ we know that $(1-r)$ will be greater than zero, which will make the numerator positive. Therefore, we can conclude that π_I will increase as p increases. We can then see that an increase in π_I will also cause our expression for l_I to increase. This increase in l_I will then in turn cause the slope of the

indifference curve to increase and the equilibrium value of σ_I to decrease. A reduction in the equilibrium value of σ_I tells us that as we increase p , jurors become less likely to vote to convict following an innocent signal. This behaviour occurs as by increasing p we are also increasing the weight that jurors place on their prior beliefs. Therefore, as p increases, jurors require further evidence in order to be convinced of a defendant's guilt.

Lastly, in order to determine the impact of r we must dive deeper as r is not only present on its own, but is also contained within the expression for l_I . To determine how an increase in r will impact the slope of the indifference curve we can differentiate the slope function with regard to r which gives:

$$\frac{d}{dr} \text{Slope}_I = \frac{(1-r)(1-l_I^2) + (2r-1) \frac{d}{dr} l_I}{(r-l_I(1-r))^2}$$

where:

$$\begin{aligned} \frac{d}{dr} \pi_I &= \frac{p(2rp+1-2r)}{(rp+(1-r)(1-p))^2} \\ \frac{d}{dr} l_I &= \frac{1}{n-1} l_I^{1-\frac{1}{n-1}} \cdot \frac{1}{c(1-\pi_I)^2} \frac{d}{dr} \pi_I. \end{aligned}$$

To begin examining this derivative, let's start by examining it piece by piece, beginning with the denominator. The denominator will always be positive due to its power. Therefore, we can focus exclusively on the numerator beginning with the terms $(1-r)(1-l_I^2)$. Given that l_t is greater than one, the term $(1-l_I^2)$ will be negative. However, the other term $(1-r)$ will be positive as long as $r < 1$. Now we can move onto examining the other terms contained within the numerator $(2r-1) \frac{d}{dr} l_I$. Intuitively, we can see that when $r > 0.5$ the term $(2r-1)$ becomes positive. The second term $(\frac{d}{dr} l_I)$ and its sign are entirely dependent on the value of $\frac{d}{dr} \pi_I$ as every other term is positive as long as l_t , c , and n are greater than zero. We can see that the denominator of $\frac{d}{dr} \pi_I$ is positive, which means we can once again focus exclusively on the numerator. By factoring p out of the numerator, we end up with the following simplified expression:

$$p(1-2r(1-p))$$

Given that $p > 0$ we can see that its sign is entirely dependent on the terms contained within the brackets. If $1-r(1-p) > 0$ we can rearrange this inequality to show that $r < \frac{1}{2(1-p)}$. This inequality tells us that when r is less than $\frac{1}{2(1-p)}$ an increase in r leads to an increase in both π_I and l_I and vice versa. Finally, putting all of this information together, we are able to deduce whether an increase in r leads to an increase or decrease in the slope of the indifference curve depends on the value of r itself. We, therefore, have to conclude that we cannot attach a definitive sign to the derivative. For instance, given the parameter values $n = 6, h = 4, c = 1, p = 0.5, \lambda = 12$ and $r = 0.1$, the equilibrium

value of σ_I is 0.872 which then increases to 0.885 when $r = 0.2$. However, if we instead set $r = 0.7$ and then increase it to 0.8, the equilibrium value of σ_I decreases from 0.25 to 0.23.

All of the above comparative statics present important implications for the set of equilibria by either increasing or decreasing the slopes of the indifference curves. This is important as by dialling up or down various parameter values, the probability that a juror votes to convict, given their signal, could change significantly. The above comparative statics and their impact on the equilibrium value of σ_I are summarised in the table below:

Parameter	Impact on σ_I
n	Increase
p	Decrease
r	Increase/Decrease
c	Increase

6 Error Probabilities

In our jury game, there are two kinds of mistakes: a Type-1 error, where an innocent person is convicted, and a Type- 2 error, where a guilty person is acquitted. Both types of errors are important to understand as they represent an outcome when the jury has failed to reach a "just" or "correct" outcome. From this point on, however, I shall focus exclusively on the Type-1 error probability as it is often considered to be the least desirable of the two. Under unanimity, the Type-1 error contour curve can be calculated by raising the probability of being pivotal when the true state is innocent to the power of $\frac{n}{n-1}$ and setting this equal to a constant to solve for σ_G . By doing this, we end up with the following, where W is a constant:

$$\hat{\sigma}_G = \frac{W^{\frac{1}{n}} - r\hat{\sigma}_I}{1 - r}$$

This equation allows us to plot the Type-1 error contour curve alongside the indifference curves for both signals in order to examine how this error relates to each signal. The Type-1 error contour curve through the non-trivial equilibrium is pictured below in Figure 2 for the unanimous case in which $r = 0.8$, $c = 1/9$ and $n = 3$. From the previous equation, we can observe that the contour curve will be linear with the slope $\frac{-r}{(1-r)}$. It is also important to note that as we move above the curve, the Type-1 error increases. Therefore, we want to be on the lowest possible curve in order to minimise the Type-1 error.

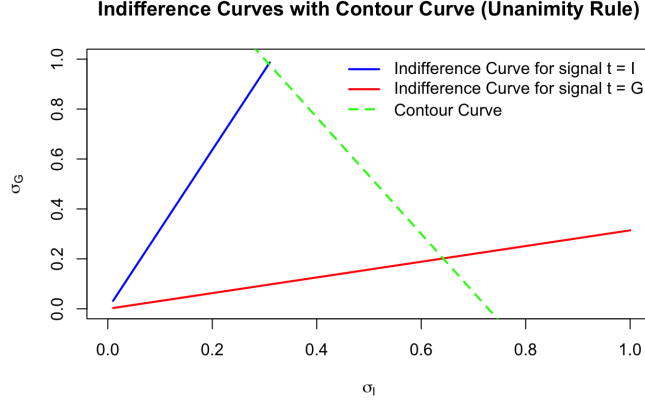


Figure 2: Indifference and contour curves under unanimity

For the case of a simple majority, this process is slightly more involved. Calculating the Type-1 error rate under a simple majority requires summing all the terms for each j in $\{h, \dots, n\}$ where j is the total number of guilty votes cast:

$$\text{Type I Error} = \sum_{k > n/2}^n \binom{n}{k} (1-p)^k p^{n-k}$$

In the above equation, we refer to p as the probability that the jury arrives at an incorrect decision in which an innocent defendant is convicted.

These equations allow us to assess how the Type-1 error rate changes as we either increase or decrease the various parameters. In the case of any non-unanimous hurdle, the construction of the curves follows a similar logic, yet they shall no longer be linear. Due to computational tedium, the visualisation of such curves has been omitted, as this does not detract from the general discussion.

6.1 The Jury Paradox

As previously discussed, the jury paradox refers to a scenario in which the probability of convicting an innocent defendant increases rather than decreases as the hurdle required for conviction increases. While the jury paradox is not confined to a single behavioural model, models that account for realistic human behaviour, such as QRE, may offer a more nuanced and accurate explanation of the phenomenon. Therefore, it is necessary to compare and contrast each model to determine which model is more adept at capturing these nuances in group decision-making.

First, we shall begin by examining the jury paradox under the Nash equilibrium framework. The jury paradox arises under Nash equilibrium as jurors vote using the same mixed strategies. Under unanimity, we know that a juror is only pivotal when everyone else has voted to convict. This pivotality condition often ensures that even jurors with strong prior beliefs vote in accordance with the will of the majority if they see or anticipate that everyone else is voting to convict. This, in turn, means that under unanimity, the Type-1 error rate is often amplified as jurors are more likely to be pivotal and are, therefore, further incentivised to vote guilty. As we decrease the hurdle required for conviction from, say, unanimity to simple majority, the likelihood that an individual juror is pivotal falls. This reduction means that jurors are more likely to vote informatively as they become less concerned with the prospect of becoming pivotal. We can, therefore, conclude that the Type-1 error rate decreases with the hurdle required for conviction under a Nash equilibrium framework.

As discussed above, the probability that any given juror becomes pivotal is largely dependent on the hurdle required for conviction. For instance, Feddersen and Pesendorfer (1998) show that under unanimity, the probability that a juror votes to convict following signal t , goes to one as n approaches infinity for all h except $h = n$. Given that $\sigma_G = 1$ in equilibrium, this result implies that the overall likelihood of conviction increases with n . The salience of this finding arises from the fact that while guilty defendants are less likely to be acquitted, innocent defendants are more likely to be convicted. As a consequence, if n is large enough, we will always observe a paradox since the Type-1 error probability is at its highest when $h = n$. However, the experimental data from Guarnaschelli, McKelvey, and Palfrey (2000) suggests that QRE may offer a way out of the paradox.

In a similar fashion to that of Feddersen and Pesendorfer (1998), tables representing juries from three to twelve members have been created in order to explore the pervasiveness of the jury paradox in the Nash equilibria where $\sigma_G = 1$. These tables can be found in Appendix A, except for the table below for a six-person jury, which was the smallest jury size to exhibit a jury paradox.

Hurdle	σ_I	Type-1 error
6	0.272	0.005
5	0.0420	0.0033
4	0	0.01696
3	0	0.098

Table 1: Type-1 error rates for different hurdles when $n = 6$.

In Table 1 (and tables in Appendix A), a jury paradox is exhibited as the hurdle for conviction moves from five to six given the parameter values of $r = 0.8$, $c = \frac{1}{9}$ and $p = 0.5$. The implications of this finding suggest that even juries with relatively few members are not immune from such a paradox which is indeed consis-

tent with what Feddersen and Pesendorfer (1998) and Guarnaschelli, McKelvey, and Palfrey (2000) find. What Table 1 adds, however, is the visualization of how quickly and to what degree a jury paradox arises. Table 1 also contributes to our understanding regarding the impact of altering various parameter values. For instance, while we already know that the Type-1 error increases with c in the unanimity case, Table 1 also shows that this is also the case when $h = 5$.

By varying the value of our c parameter, we can investigate to what degree risk aversion impacts the existence and impact of the jury paradox. This can be done by repeating the above process of producing tables representing juries from three to twelve members, with the added alteration of increasing c from $\frac{1}{9}$ to $\frac{2}{9}$. Having completed this process, we are able to see that by increasing c the smallest jury size to exhibit a paradox remains at six jurors. All of these tables can be located in Appendix B, with the exception of the table representative of six jurors, which is presented below.

Hurdle	σ_I	Type-1 error
6	0.362	0.0143
5	0.102	0.00814
4	0	0.01696

Table 2: Type-1 error rates for $n = 6$ and $c = \frac{2}{9}$.

In Table 2, we can see that while the smallest jury size to exhibit a paradox remains at six jurors, the Type-1 error probabilities and the σ_I values have significantly changed. This further proves our previous comparative static analysis. As shown in Table 1, Table 2 also exhibits a jury paradox as we increase c for $h = 5$. In particular, we can see that as the hurdle required for conviction increases from five to six, the Type-1 error probability increases from 0.00814 to 0.0143. The fact that the Type-1 error probabilities and the σ_I values have significantly changed reveals important information about jurors' attitudes towards risk.

The increase in the value of c causes the payoff associated with acquitting a guilty defendant to become more negative. This means that as we increase the value of c , the payoff associated with acquitting a guilty defendant approaches that received from convicting an innocent defendant (-1). As this occurs, jurors slowly become more and more indifferent towards convicting an innocent defendant or acquitting a guilty defendant. This alteration in risk attitude is clearly evidenced in Table 2, as the probability that a juror votes to convict upon observing an innocent signal has sharply increased from the values shown in Table 1.

We not only know, but have also shown that by increasing the value of our c parameter, the Type-1 error probability increases as well under unanimity. It

is therefore worth investigating whether this effect also persists for $h = 5$ across the full range of potential c values and how it compares with the Type-1 values for $h = 6$.

The table below (Table 3) sets p , n , and r at the usual values of 0.5, 6, and 0.7, respectively, while showing how the Type-1 error rate changes across the hurdles in $\{5, 6\}$ as c increases.

c	$h = 5$ (Type -1) error	$h = 6$ (Type -1) error
0.1	0.010935	0.003761578
0.2	0.01235195	0.01131895
0.3	0.02214897	0.02210743
0.4	0.03364418	0.03601763
0.5	0.04661943	0.05303657
0.6	0.06092021	0.07319213
0.7	0.07642623	0.0965339
0.8	0.09303788	0.1231247
0.9	0.1106691	0.1530361
1.0	0.1292431	0.1863461

Table 3: Type -1 error rates for $n = 6$ and $h = \{5, 6\}$

Table 3 suggests that at least for the given parameter values, we can extend our finding which states that under unanimity the Type-1 error rate increase with c , to state that the same can be said when $h = 5$. However, more importantly, the table shows us this only occurs for c values above 0.3. This implies that, at least for these parameter values, c only needs to be increased slightly in order to make jurors more likely to convict.

7 Quantal Response Equilibrium

In this paper, we shall examine the properties and results of another model which, unlike Nash equilibrium, incorporates probabilistic choice. A quantal response function maps expected payoffs to choice probabilities where one is an increasing function of the other (Goeree, Holt, and Palfrey 2020). A quantal response equilibrium is then a fixed point of the quantal response functions in the same way that Nash Equilibrium is a fixed point of best responses in standard game theory (Goeree, Holt, and Palfrey 2020). It is important to note that at a quantal response equilibrium, the distribution of choice probabilities will also match the distribution of a player's belief distributions regarding the actions that other players are likely to make (Goeree, Holt, and Palfrey 2008).

In quantal response equilibria, the logit, and probit models are standard options used to model quantal response functions. While all three models have

their place in the analysis of various problems, the logit model is the most commonly used of the three. The popularity and widespread use of the logit model is most likely due to the fact that it not only allows for the use of negative payoffs but also allows responses to be expressed in relation to a responsiveness or "noise" parameter (λ) that can be used to scale different levels of rationality or precision Guarnaschelli, McKelvey, and Palfrey 2000. A greater lambda means that actions which result in greater payoffs are played with greater probability Goeree, Holt, and Palfrey 2020. As lambda approaches infinity and players become more rational through the reduction of noise, the quantal response equilibrium will converge to the corresponding Nash equilibrium. This is due to the fact that as we increase lambda, the probability that jurors vote in a payoff-maximizing manner becomes close or equal to one. This is shown in the mathematical form of the logit quantal response function below :

$$\sigma_t = \frac{\exp(\lambda EU(t, G))}{\exp(\lambda EU(t, G)) + \exp(\lambda EU(t, I))}$$

and

$$\lim_{\lambda \rightarrow \infty} \frac{e^{\lambda EU(t, G)}}{e^{\lambda EU(t, G)} + e^{\lambda EU(t, I)}} = \begin{cases} 1, & \text{if } EU(t, G) > EU(t, I), \\ 0, & \text{if } EU(t, G) < EU(t, I). \end{cases}$$

The above limit shows that as λ approaches zero the exponent on e also approaches zero for every action. This means that the differences in expected utilities from each action slowly reduce until every action is chosen with equal probability given that $e^0 = 1$ and $\frac{1}{1+1} = \frac{1}{2}$. We can also observe that in a two-action game, if both actions have the same expected utilities, then both will be played with probability 0.5 irrespective of lambda. Due to the complexity of these calculations, solutions require numerical methods to be reached. For instance, Guarnaschelli, McKelvey, and Palfrey (2000) shows in Figure 4 that lambda is non-monotonic in the symmetric part of the QRE correspondence. This produces a multiplicity problem with the QRE correspondence that causes the QRE to double back on itself and become multi-valued around the point where $\lambda = 15$. For more information, the reader is directed to the paper by Guarnaschelli, McKelvey, and Palfrey (2000).

7.1 Quantal Response Equilibrium in this model

In order to incorporate QRE into our existing model, we must find the necessary conditions for an arbitrary quantal response function. There are two key features of the quantal response function that we shall use that must be emphasised. Before emphasising these features, it is important to remember that we are referring to a two-action game, given that jurors only have the option to vote to either convict or acquit. The first feature worth emphasising is that the action with the higher expected utility will be played with probability > 0.5 as shown in many papers, including that from Goeree, Holt, and Palfrey (2020). The second key feature is that if both actions share the same expected utility,

then they will both be played with probability 0.5 which is also a key feature discussed in many papers, including that from McKelvey and Palfrey (1995).

To visualise how QRE works within our model, we can further utilise Figure 1. We can cut Figure 1 into four equal quadrants at the point $(0.5, 0.5)$ as the probabilities of both actions are equal to 0.5 when the expected payoff difference between the two is zero (Goeree, Holt, and Palfrey 2020). For any σ_I or σ_G value that is greater than 0.5 the QRE feasibility region will lie above the curve, and for any value below 0.5, the feasibility region will lie below the curve. Let's begin by looking at the blue curve for $t = I$.

First, consider a juror who is responding to a strategy profile above the blue indifference curve. At any point above the curve, we already know that regardless of the signal, it is optimal to vote to convict rather than acquit. This means that a quantal best response must play $v = G$ with probability > 0.5 and hence be located in the upper right quadrant. However, no such points are located in this region, meaning that there can't be a QRE above the blue indifference curve.

Now consider a juror who is responding to a strategy profile located below the red indifference curve. At any point below the curve, we already know that it is optimal to vote to acquit, regardless of which signal is observed. Therefore, a quantal best response must play $v = G$ with probability < 0.5 and hence be located in the bottom left quadrant. We can see that there is indeed an array of points below the red curve in this quadrant. We can, therefore, shade this section orange to confirm that a QRE can lie within this region.

Next, consider a juror who is responding to a strategy profile located below the blue indifference curve and above the red indifference curve. Given that this strategy is located between the indifference curves, we know that the best response is to vote to convict following a guilty signal and to acquit following an innocent signal. Therefore, a quantal best response must set $\sigma_G > 0.5$ and $\sigma_I < 0.5$. These conditions mean that a QRE can exist inside the green shaded region.

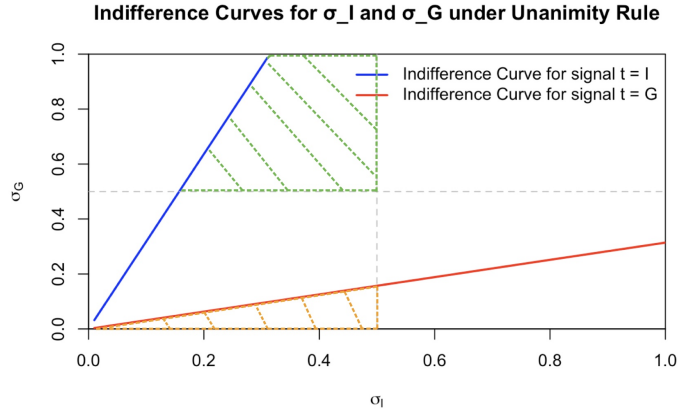


Figure 3: QRE feasibility regions $n = 3$

As discussed previously, as we alter various parameter values, the curves can become either flatter or steeper. This, of course, will impact the size and location of the QRE feasibility region. For example, consider Figure 4 which uses the same parameter values as Figure 3 apart from n , which has increased to six. We can see straight away that Figure 4 has an additional region in which a QRE can live. This extra green region exists due to the fact that the blue curve has been shifted clockwise to the point where there is now a region which facilitates a best response of $v = G$ when responding to a point above the curves. We can also see that by increasing n , the region below the red curve has increased in area while the region between the curves has shrunk.




Figure 4: QRE feasibility regions $n = 6$

By looking at Figures 2 and 3, we can conclude that, in principle, QRE has

the potential to either raise or lower the Type-1 error probability relative to the Nash equilibrium. However, in the case of Figure 4, QRE will lower the Type-1 error probability due to the reduced area of the regions where the QRE can lie.

8 Logit QRE

As previously discussed, the logit model is one of the most commonly used models to configure quantal response functions. The main rationale behind using the logit model stems from the fact that it satisfies symmetry, handles both negative and positive payoffs, and is strictly positive for all actions Goeree, Holt, and Palfrey (2020). To incorporate the logit model, we must first formulate the logit quantal response function in relation to the jury game:

$$\sigma_t = \frac{\exp(\lambda EU(t, G))}{\exp(\lambda EU(t, G)) + \exp(\lambda EU(t, I))}$$

where σ_t is the probability that a juror votes to convict upon receiving signal t . The expected utility received from voting to either convict or acquit are denoted as $EU(t, G)$ and $EU(t, I)$ respectively. Finally, we have the lambda parameter, which, as we know, represents the degree to which jurors behave rationally.

As discussed in the previous section on the general framework of QRE, we must solve this logit quantal response function for a fixed point in order to establish a QRE. In order to do this, we must employ one of the numerical methods mentioned previously. In this case, I have chosen to use a fixed-point optimization method within R Studio that utilises a non-linear least squares approach. One of the reasons for choosing this method over others, such as standard fixed point iteration, is that provides internal consistency through its reduced reliance on explicit starting points.

9 QRE and The Jury Paradox

Given that we have touched on the jury paradox through the lens of Nash equilibria, we can move on to examining the paradox through the lens of quantal response equilibrium (QRE). As we know, QRE introduces stochastic choice into decision-making via the introduction of a noise parameter (λ). By scaling up and down this noise parameter, we can explore whether the shift from deterministic to probabilistic choice either increases or decreases the prevalence and pervasiveness of the jury paradox. As lower values of lambda correspond to a greater degree of stochastic behaviour, a jury paradox is more likely to occur as lambda approaches zero rather than infinity. This is proven by the empirical findings from Guarnaschelli, McKelvey, and Palfrey (2000) and is consistent

with their theorems regarding the logit equilibrium. In theorem one, they show that even under unanimity, as n increases, the probability of conviction goes to zero. Given that the probability of conviction approaches zero as n increases, theorem one also implies that the probability of acquittal approaches one as n increases.

Theorem two supports the pervasiveness of theorem one by showing that even with the introduction of a straw poll, the findings of theorem one remain the same. Both of these theorems detail just how QRE can better manage and account for the jury paradox compared to Nash equilibrium predictions. However, it is important to note that while the probability of conviction approaches zero as n increases, guilty defendants are more likely to be acquitted than convicted. Yet, as long as the payoff associated with ensuring the freedom of innocent defendants is greater than that for convicting guilty defendants, jurors will be satisfied with this arrangement.

Now that we have discussed how QRE addresses the jury paradox under unanimity, it is important to explore its impact under other hurdles, such as simple majority. Conveniently, Guarnaschelli, McKelvey, and Palfrey (2000) does this with reference back to theorem one and two. In particular, Guarnaschelli, McKelvey, and Palfrey (2000) details that under a simple majority rule for conviction, the preference for acquittal, which is exhibited under unanimity, is no longer present. However, the disappearance of this preference comes at a cost. Guarnaschelli, McKelvey, and Palfrey (2000) states that this cost is equal to a fifty per cent increase in the rates of incorrectly convicting innocent defendants. Therefore, QRE offers the most protection to innocent defendants when a unanimous verdict is required.

10 Results and Discussion

In order to explore and visualise how the logit QRE model can account for the jury paradox within the jury model, three tables have been created. Table 4 shows the values of r and c that correspond to a jury paradox under the Nash equilibrium model. Table 4 represents the case where $n = 6$ and $h = \{4, 5, 6\}$ as this is the lowest jury size to exhibit the paradox given the parameter values underpinning Tables 1 and 2. The first step in creating the tables is to search for equilibria of the form $(\sigma_I, 1)$ and reject cases where this form is not met. The cells containing -1 represent the case where the required equilibrium structure is not met for all hurdle values, such that σ_I isn't within the range of $[0, 1]$. Next, each relevant cell contains a set of three numbers such as $(0, 1, 0)$. The first number indicates whether a jury paradox arises as we move from $h = 4$ to $h = 5$. If a jury paradox does arise, then this place will be filled with a 1 and if not, it will be filled with a 0. The same logic also applies to the second number which represents whether a jury paradox arises as we move from $h = 5$ to $h = 6$. Finally, the third number indicates whether a jury paradox arises as

we go from $h = 4$ directly to $h = 6$

$c \backslash r$	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95
0.1	-1	-1	0,0,0	0,0,0	0,0,0	0,1,0	0,0,0	0,1,1	1,1,1
0.2	-1	0,0,0	0,0,0	0,0,0	0,1,0	0,1,0	0,1,1	1,1,1	1,1,1
0.3	0,0,0	0,0,0	0,0,0	0,0,0	0,1,0	0,1,1	1,1,1	1,1,1	1,1,1
0.4	0,0,0	0,0,0	0,0,0	0,1,0	0,1,0	1,1,1	1,1,1	1,1,1	1,1,1
0.5	0,0,0	0,0,0	0,0,0	0,1,0	0,1,1	1,1,1	1,1,1	1,1,1	1,1,1
0.6	0,0,0	0,0,0	0,0,0	0,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1
0.7	0,0,0	0,0,0	0,1,0	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1
0.8	0,0,0	0,0,0	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1
0.9	0,0,0	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1
1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1	1,1,1

Table 4: NE (c, r) pairs for $n = 6$

Table 4 tells us that under a Nash equilibrium framework, the jury paradox thrives when both r and c are large. However, Table 4 shows us that even for relatively low values of r such as 0.6, as long as c is relatively large, a jury paradox still prevails. This suggests that c has a greater impact in regards to determining the behaviour of jurors compared to r . This finding indicates that jurors are highly risk-averse and prefer to convict even when they are largely uncertain about the accuracy of their signals. This finding also suggests that the jury’s threshold of reasonable doubt is relatively low, meaning that they don’t require much to be swayed into voting to convict.

The second table (Table 5) is constructed in a similar way to that of Table 4, yet it incorporates QRE instead of a Nash equilibrium framework. Switching from a Nash equilibrium framework to a QRE framework requires the introduction of a noise parameter (λ). Table 5 sets $\lambda = 10$ as this allows us to avoid the multiplicity issue found by Guarnaschelli, McKelvey, and Palfrey (2000) in Figure 4. This ensures that the equilibrium Type-1 error is unambiguous. It is important to note, however, that the tables created in this paper explore a range of parameter values for which our computations showed no evidence of multiple equilibria at any valid parameter configuration.

Table 5 copies over the -1 values from Table 4 and then assigns the relevant 0 and 1 markings as determined by the logit quantal response function.

$c \backslash r$	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95
0.1	-1	-1	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.2	-1	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.3	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.4	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.5	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.6	0,0,0	0,0,0	0,0,0	0,1,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.7	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.8	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.9	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
1	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0

Table 5: QRE (c, r) pairs for $\lambda = 10$

Simply by glancing at Table 5, we can immediately determine that the introduction of QRE has greatly managed and accounted for the jury paradox exhibited in Table 4. To put this into perspective, the one hundred and sixty-three instances of the jury paradox exhibited in Table 4 have shrunk to one by simply switching from a Nash equilibrium framework to a QRE one. This finding implies that even for relatively low values of lambda, QRE still does a good job at managing the jury paradox.

Finally, a third table (Table 6) has been created in order to capture the impact of increasing lambda from 10 to 25. Table 6 paints an interesting picture as it reveals that even increasing lambda by fifteen doesn't hugely impact the degree to which the QRE is able to manage the paradox. In saying this however, we can see that Table 6 does exhibit more jury paradoxes than Table 5 for higher c and r values. This suggests that while setting $\lambda = 25$ still manages the paradox well, the responses are closer to those seen under the Nash equilibrium model in Table 4.

$c \backslash r$	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95
0.1	-1	-1	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.2	-1	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.3	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.4	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0
0.5	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	1,0,1
0.6	0,0,0	0,0,0	0,0,0	0,1,0	0,0,0	0,0,0	0,0,0	0,0,0	1,0,1
0.7	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	1,0,1	1,0,1
0.8	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	1,0,1	1,0,1	1,0,1
0.9	0,0,0	0,0,0	0,1,0	0,0,0	0,0,0	1,0,0	1,0,0	1,0,1	1,0,1
1	0,0,0	0,0,0	0,0,0	0,0,0	1,0,0	1,0,0	1,0,1	1,0,1	1,0,1

Table 6: QRE (c, r) pairs for $\lambda = 25$

Both Tables 5 and 6 are worth comparing with the experimental findings from Guarnaschelli, McKelvey, and Palfrey (2000). For instance, the experimental behaviour they present does not exhibit a paradox which is a direct contradiction to what Tables 5 and 6 uncover. This is most likely due to the fact that Guarnaschelli, McKelvey, and Palfrey (2000) varies lambda to fit their data, allowing for different lambda values under both unanimity and simple majority. Therefore, it is not surprising that their QRE values don't generate a paradox. However, this means that their results don't necessarily offer the best comparison between Nash equilibrium and QRE. Instead, Tables 4, 5, and 6 offer a more balanced comparison while also showing that even if we require the same lambda for all voting rules, QRE still does very well at eliminating the paradox. We can see this fact further visualized in Figure 5 as both QRE curves are significantly different to that for the Nash equilibrium.

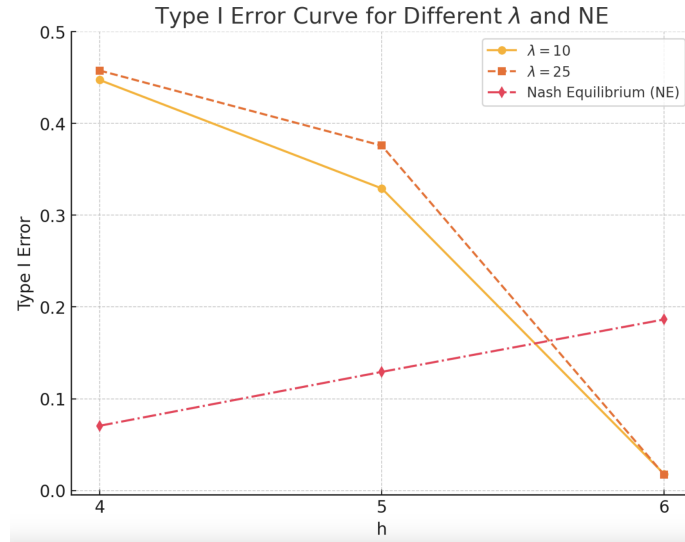


Figure 5: Type-1 error curve for different λ and NE

11 Significance and Limitations

The above integration and deployment of QRE in the context of the proposed jury game has confirmed that QRE does indeed provide a more suitable and realistic analysis of the data. Specifically, QRE provides a more pragmatic analysis of individual and collective behaviour compared to traditional game theoretic models such as Nash equilibrium. The real significance of this in the context of the jury game, however, relates to how the jury paradox exhibits itself across various parameter values. For instance, Tables 5 and 6 detail the significance of asymmetric payoffs and their impact on the pervasiveness of the paradox.

In particular, these tables suggest that the paradox exhibited under the Nash equilibrium relies on high c and r values. Unsurprisingly, the ability of QRE to offset the paradox relies on λ being fairly low. The findings in this paper detail just how important it is to update traditional game theory models in order to provide a clearer picture of realistic decision-making in an applied context.

While this paper has identified that QRE does indeed offer a more robust and realistic assessment of predicted behaviour, it is important to note that it does have limitations. For instance, it is relatively computationally intensive, which can make it challenging to model games with more than two actions. Also, as mentioned previously, QRE can often encounter a multiplicity issue for certain λ values. This can make it difficult to ascertain a unique equilibrium solution for certain values.

Although this paper does indeed provide new insights that support the use of stochastic rather than deterministic behavioural frameworks, it must be acknowledged that there is more work to be done. For instance, in order to provide a more thorough comparison between QRE and Nash equilibria frameworks, further research may need to consider altering the foundations and assumptions of the model presented by Feddersen and Pesendorfer (1998). For instance, the research by Anderson et al. (2015), Coughlan (2000), and Eyster and Rabin (2005) suggests that altering some of the model's parameters and removing some of its strict simplifying assumptions may have a significant impact on the pervasiveness of the jury paradox under a Nash equilibrium framework.

In reference to the QRE model, further research can and should be done in order to further flesh out the robustness of the model. For example, in most studies, including this one, λ is generally either chosen arbitrarily or selected in order to provide an optimal match to experimental data. If future research is able to provide a more sound rationale regarding the selection and implications of this parameter, a more cohesive and cemented logic can be applied across the relevant theoretical and experimental spaces.

12 Conclusion

The field of game theory has evolved immensely since its conception and continues to do so to this day. From Condorcet's prominent paper published back in 1785 to the more recent work on stochastic choice models, game theory continues to be rooted in the discovery, predictability, and rationality of human behaviour. I believe this study has contributed to and assessed the importance of both older and more modern approaches to modelling decision-making while acknowledging the importance of each.

This study set out to examine the predictive and applied accuracy of Quantal Response Equilibrium (QRE) compared to Nash Equilibrium (NE) in voting

games, motivated by the observation that real-world decision-making often deviates from the strict rationality assumptions underlying traditional game theory models. While NE assumes that players are entirely rational in the sense that they make choices without error, empirical evidence suggests that factors such as bounded rationality and decision noise frequently lead to deviations from Nash equilibrium behaviour. QRE introduces a probabilistic decision-making framework, allowing for more flexible and realistic predictions of strategic behaviour under uncertainty.

To explore this, we modelled a jury voting game in which individuals receive private signals and vote strategically based on their expected payoffs. Using computational simulations, we analysed how equilibrium predictions shift under varying levels of expected payoffs, signal precision, strategic uncertainty and others. The findings indicate that as uncertainty increases, QRE better captures observed deviations from NE, suggesting that real-world jurors may not always play best responses deterministically but rather choose probabilistically based on a combination of expected utility and expectations regarding the behaviour and beliefs of others.

Overall, this study highlights the need for decision models that better account for authentic, real-world decision-making. As decision-making grows increasingly complex—whether in jury voting or political settings—models like QRE, which integrate bounded rationality and probabilistic choice, will be essential for understanding and predicting real-world strategic behaviour.

References

- Anderson, Lisa R. et al. (2015). “An Experimental Study of Jury Voting Behavior”. In: *The Political Economy of Governance*. Ed. by Norman Schofield and Gonzalo Caballero. Series Title: Studies in Political Economy. Cham: Springer International Publishing, pp. 157–178. ISBN: 978-3-319-15550-0 978-3-319-15551-7. DOI: 10.1007/978-3-319-15551-7_8.
- Austen-Smith, David and Jeffrey S. Banks (1996). “Information Aggregation, Rationality, and the Condorcet Jury Theorem”. In: *The American Political Science Review* 90.1. Publisher: [American Political Science Association, Cambridge University Press], pp. 34–45. ISSN: 0003-0554. DOI: 10.2307/2082796.
- Bouton, Laurent, Aniol Llorente-Saguer, and Frédéric Malherbe (Mar. 2017). “Unanimous rules in the laboratory”. In: *Games and Economic Behavior* 102, pp. 179–198. ISSN: 08998256. DOI: 10.1016/j.geb.2016.12.001.
- Condorcet, Jean Antoine Nicolas de Caritat Mis de (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Google-Books-ID: tEWIepFiLXQC. Imprimerie royale. 506 pp.
- Coughlan, Peter J. (June 2000). “In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting”. In: *American Political Science Review* 94.2, pp. 375–393. ISSN: 0003-0554, 1537-5943. DOI: 10.2307/2586018.
- Eyster, Erik and Matthew Rabin (Sept. 2005). “Cursed Equilibrium”. In: *Econometrica* 73.5, pp. 1623–1672. ISSN: 0012-9682, 1468-0262. DOI: 10.1111/j.1468-0262.2005.00631.x.
- Feddersen, Timothy and Wolfgang Pesendorfer (1998). “Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting”. In: *The American Political Science Review* 92.1. Publisher: [American Political Science Association, Cambridge University Press], pp. 23–35. ISSN: 0003-0554. DOI: 10.2307/2585926.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2008). “Quantal response equilibria”. In: Publisher: Palgrave Macmillan. DOI: 10.5167/UZH-58312.
- (Oct. 15, 2020). “Stochastic game theory for social science: a primer on quantal response equilibrium”. In: *Handbook of Experimental Game Theory*. Ed. by C. M. Capra et al. Edward Elgar Publishing. ISBN: 978-1-78536-333-7 978-1-78536-332-0. DOI: 10.4337/9781785363337.00008.
- Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey (2000). “An Experimental Study of Jury Decision Rules”. In: *The American Political Science Review* 94.2. Publisher: [American Political Science Association, Cambridge University Press], pp. 407–423. ISSN: 0003-0554. DOI: 10.2307/2586020.
- McKelvey, Richard D. and Thomas R. Palfrey (July 1995). “Quantal Response Equilibria for Normal Form Games”. In: *Games and Economic Behavior* 10.1, pp. 6–38. ISSN: 08998256. DOI: 10.1006/game.1995.1023.

13 Appendix A: Jury Paradox Comparison When

$$c = \frac{1}{9}$$

The below tables incorporate the usual parameter values where $p = 0.5, r = 0.8$ and $c = \frac{1}{9}$

Hurdle	σ_I	Type-1 error
3	0	0.008
2	0	0.104

Table 7: Table for $n = 3$

Hurdle	σ_I	Type-1 error
4	0.057	0.00365
3	0	0.0272
2	0	0.1808

Table 8: Table for $n = 4$

Hurdle	σ_I	Type-1 error
5	0.176	0.005
4	0	0.0067
3	0	0.0579

Table 9: Table for $n = 5$

Hurdle	σ_I	Type-1 error
6	0.272	0.005
5	0.0420	0.0033
4	0	0.01696

Table 10: Table for $n = 6$

Hurdle	σ_I	Type-1 error
7	0.348	0.00575
6	0.132	0.00423
5	0	0.00467
4	0	0.033

Table 11: Table for $n = 7$

Hurdle	σ_I	Type-1 error
8	0.411	0.00611
7	0.21	0.00494
6	0.033	0.00244
5	0	0.010

Table 12: Table for $n = 8$

Hurdle	σ_I	Type-1 error
9	0.463	0.0064
8	0.2755	0.0055
7	0.106	0.0031
6	0	0.0031
5	0	0.01958

Table 13: Table for $n = 9$

Hurdle	σ_I	Type-1 error
10	0.5066	0.0066
9	0.332	0.00596
8	0.171	0.00364
7	0.0270	0.00165
6	0	0.006

Table 14: Table for $n = 10$

Hurdle	σ_I	Type-1 error
11	0.544	0.0068
10	0.380	0.0063
9	0.2275	0.0041
8	0.088	0.0021
7	0	0.0019
6	0	0.01165

Table 15: Table for $n = 11$

Hurdle	σ_I	Type-1 error
12	0.5750	0.0069
11	0.423	0.00666
10	0.277	0.0045
9	0.144	0.0025
8	0.023	0.0011
7	0	0.004

Table 16: Table for $n = 12$

14 Appendix B: Jury Paradox Comparison When $c = \frac{2}{9}$

The below tables incorporate the usual parameter values where $p = 0.5$ and $r = 0.8$.

Hurdle	σ_I	Type-1 error
4	0.145	0.01
3	0	0.0272

Table 17: $n = 4$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
5	0.2680	0.012
4	0	0.0067
3	0	0.0579

Table 18: $n = 5$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
6	0.362	0.0143
5	0.102	0.00814
4	0	0.01696

Table 19: $n = 6$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
7	0.435	0.015
6	0.2	0.01
5	0	0.00467
4	0	0.033

Table 20: $n = 7$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
8	0.5	0.0156
7	0.274	0.012
6	0.0785	0.0055
5	0	0.010

Table 21: $n = 8$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
9	0.541	0.0163
8	0.34	0.013
7	0.154	0.0069
6	0	0.0031
5	0	0.01958

Table 22: $n = 9$, $c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
10	0.581	0.0168
9	0.4	0.0136
8	0.22	0.01
7	0.064	0.004
6	0	0.006

Table 23: $n = 10, c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
11	0.614	0.0172
10	0.44	0.0144
9	0.3	0.01
8	0.13	0.0044
7	0	0.0019
6	0	0.01165

Table 24: $n = 11, c = \frac{2}{9}$

Hurdle	σ_I	Type-1 error
12	0.579	0.02
11	0.432	0.02
10	0.288	0.019
9	0.15	0.016
8	0.0145	0.0118
7	0	0.04

Table 25: $n = 12, c = \frac{2}{9}$