

DO JUDGES ANCHOR TO LITIGANTS' NUMBERS?

Anthony Niblett
University of Toronto Faculty of Law

Very preliminary draft. Preliminary results only.
Please do not cite or circulate without permission. Thank you!

July 10, 2022

1. INTRODUCTION

Suppose there is a dispute where the only issue is the amount of damages that the defendant should pay the plaintiff. The plaintiff asks for a high number. The defendant puts forward a low number. Do these numbers put forward by the litigants influence the judge's decision? If so, how exactly does the judge use the litigants' numbers? These are the questions I explore in this paper.

One can put forward some general hypotheses of how judges treat the litigants' numbers:

- (1) **The judge ignores the litigants' numbers.** Perhaps the judge feels that the litigants are biased in their views. The numbers offer more noise than signal value. The judge independently decides what the correct damages award should be.
- (2) **The judge may split the difference between the plaintiff's high number and the defendant's low number.** While the biases of the litigants' numbers would offset, if parties know that judges split the difference, it could generate incentives for the parties to exaggerate their position.
- (3) **The judge awards the 'most reasonable' or 'fairest' number put forward by the two litigants.** If parties know that judges behave in this way, then they are incentivized to put forward a number they feel is reasonable.

The way in which judges use the litigants' numbers is important because it will influence litigation strategy. By extension, it impacts upon the likelihood of settlement.

Suppose judges always split the difference. As noted above, *if* the litigants know that the judges simply split the difference, this generates incentives for parties to exaggerate their position. Plaintiffs would ask for unreasonably high amounts; defendants offer unreasonably low amounts. Thus, if judges split the difference, it would likely exacerbate the disagreement between the two parties.

On the other hand, suppose judges award damages equal to the number put forward by the most reasonable party. If the parties know this, then they each have an incentive to gravitate towards a reasonable claim or offer. This has the effect of reducing the area of disagreement between the two parties and is more likely to result in pre-litigation settlement.

This paper is empirical and descriptive. I look at the awards of damages made by judges and those figures to the numbers put forward by the plaintiff and the defendant. I create a novel dataset of cases in Canadian employment law where damages are at issue. The cases all deal with the same issue: payment in lieu of reasonable notice. Under Canadian common law, if a worker is dismissed without cause, the worker is entitled to a reasonable notice period, or payment in lieu of the notice period. What is 'reasonable' in any case, however, depends on the facts. It is a frequent source of disagreement between plaintiff employees and defendant employers. There has been a wealth of litigation of this issue in Canada, with (literally) thousands of cases decided by courts over the past 50 years.

I have granular data from 1,757 judicial decisions over the years 1997 to 2019. In 601 of these cases, I know precisely what numbers were proposed by both the plaintiff employee and the defendant employer because the judge includes the litigants' numbers in the written opinion. In these cases, I know how much the plaintiff employee asks for and I know how much the defendant employer is willing to offer. I define the "area of disagreement" as the difference between the high number sought by the plaintiff employer and the low number offered by the defendant employer.

Figure 1 previews the results. **Figure 1** shows a histogram of the outcomes of the 601 cases. It presents the damages awarded by the judge, normalizing the area of disagreement to $[0,1]$ space. Here, 0 represents that the judge awarded damages equal to a low number offered by the defendant employer, 1 represents that the judge awarded damages equal to a high number claimed

by the plaintiff employee, and 0.5 represents an outcome that splits the difference between the two.

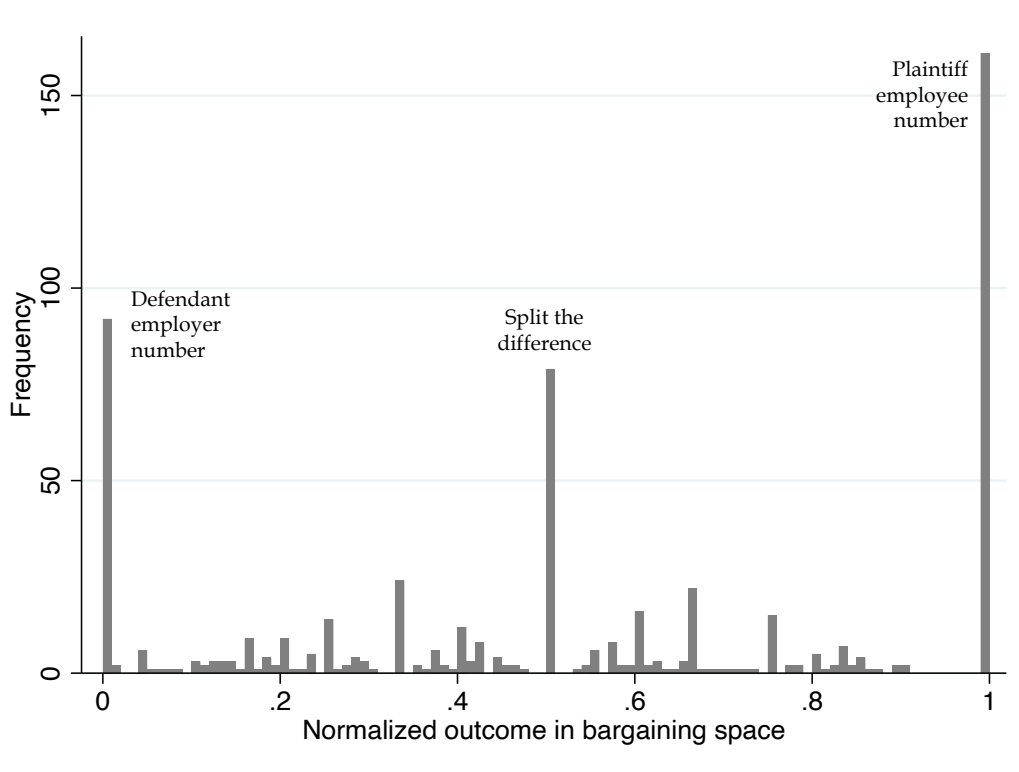


Figure 1: Frequency of normalized outcomes, locating the damages awarded by the judge within the *area of disagreement*. A normalized outcome of 0 means that the judge awarded the low number put forward by the defendant. A normalized outcome of 1 means that the judge awarded the high number put forward by the plaintiff. A normalized outcome of 0.5 means that the judge split the difference.

To briefly summarize the evidence from **Figure 1**:

- Judges do take the litigant numbers into account. Judges appear to anchor, at least in part, to the litigant numbers.
- In about 42% of the 601 cases, the judge awards either the high number put forward by the plaintiff or the low number put forward by the defendant. The modal outcome is the high number sought after by the plaintiff.
- Judges split the difference between the plaintiff’s high number and the defendant’s low number in around 13% of cases.

The simple graphical depiction in Figure 1 is illuminating. But more is needed to test the hypotheses. I have extremely granular data about the facts of the cases. With these data, I predict the likely strength of the plaintiff employee’s

case. That is, I predict a likely outcome (*i.e.*, the likely reasonable notice) simply from the facts of the case. This allows me to calculate how much of a difference is made from additionally knowing the plaintiff's high number and the defendant's low number. These numbers matter. The litigants' numbers help explain about 29% of the unexplained variation in the judicial outcomes when only the facts of the case are taken into account. Including litigants' numbers in the model also increases the accuracy of outcome predictions by about 17%.

There are two ways I measure the reasonableness or fairness of the litigants' numbers. *First*, I compare the litigants' numbers to the predicted outcome. When the predicted outcome is closer to the plaintiff's high number, judges are more likely to gravitate towards the plaintiff's number. And when the predicted outcome is closer to the defendant's low number, judges are more likely to award this low number. But the evidence that they gravitate towards the most reasonable outcome - by this measure - is not overwhelming.

Second, I compare the numbers asked for a plaintiff to those asked for by other plaintiffs - and similarly, the numbers offered by a defendant to those asked for by other defendants. Using the granular data on facts, I make predictions about how the litigants would be expected to behave.

The plaintiffs are ranked in terms of how relatively aggressive their claim is. There is a strong (negative) relationship between how aggressive - and potentially unreasonable - the claim is and the likelihood of the judge anchoring to the plaintiffs' claim. In the 100 cases where the plaintiff acts most passively, judges award the plaintiffs' claim in about half the cases. In the 100 cases where the plaintiffs act the most aggressively, however, the judge awards the claim in 3 cases. When a plaintiff asks for a number than it is significantly higher than expected, the judge is far less likely to anchor to the plaintiff's number. Interestingly, though, there is no statistical relationship between the aggressiveness of the plaintiff's claim and the likelihood that the judge will gravitate to the defendant's number.

I also rank the defendants by the relative aggressiveness of their offer. There is a strong negative relationship between anchoring to the defendants' offer and the aggressiveness of the offer. As the offer becomes more generous, the likelihood that the judge will anchor to that number increases. In the 100 cases where the defendant is most aggressive, the judge awards the number put forward by the defendant in 2 cases. In the 100 cases where the defendant is least aggressive, the judge anchors to their number in nearly half the cases. There appears to be no statistical relationship between the aggressiveness of the defendant's offer and the likelihood that the judge will anchor to the plaintiff's number.

* * *

To sum up, there is strong evidence that judges anchor toward one of the litigants' numbers. Judges are more likely to gravitate toward the number put forward by a party if that number is more reasonable. Judges anchor on plaintiff claims when plaintiffs are passive; judges anchor to defendant offers when defendants are generous.

The remainder of the paper is set out as follows. Part 2 presents a literature review (incomplete.) Part 3 describes the legal setting and set out the hypotheses. Part 4 describes the data. Part 5 introduces preliminary results. Part 6 tests the hypotheses. *[Discussion and conclusion to be written.]*

2. LITERATURE REVIEW

There is a robust literature demonstrating how decision makers are influenced by numbers to which they were previously exposed. Tversky and Kahneman (1974) constructed a classical experiment. They showed that participants' estimates of various quantities, such as the number of African countries, were influenced by exposure to random numbers. Participants given a higher number answered with a higher estimate. These authors theorized that the participants made insufficient adjustments to the starting value in arriving at their estimate.

Anchors influence the decisions of judges and juries. Rachlinski, Wistrich, & Guthrie (2015) presented judges with hypothetical cases and found that damages awards and criminal sentences were affected by misleading numeric reference points. A prosecutor's sentencing demand may serve as an anchor, and thus disproportionate weight may be given to the prosecutor's demand, recommendation, or sentencing guidelines (*e.g.*, Englich 2006; Bennett 2014; Xifen & Ma 2018). Similarly, the recommendations of lawyers in civil claims for damages has been shown to influence juries (King 2003; Hans & Reyna 2011; Diamond, Rose, Murphy, Meixner 2011.) Caps on punitive damage awards have been found to engender anchoring effects in punitive and compensatory damage award decisions (Robbennolt & Studebaker 1999).

Chapman & Johnson (1999) show that the anchoring effect may result from the selective activation of information that is consistent with the anchor. To illustrate this idea, judges presumably consider the anchor to be a plausible answer and search to confirm whether this hypothesis is correct. In doing so, judges search for ways that their answer is similar to the anchor and thus activate features of the target similar to the anchor.

Where advocates from both sides propose numbers, each may exert influence to create a “combined anchoring” effect. Stein and Drouin (2018) found that defense attorneys may reduce the anchoring effect by countering the prosecution’s anchor. Imhoff and Nickolaus (2021) observed a strong correspondence between the distance between the demands of both parties and the dispersion of sentencing decisions. Decision-makers did not seem to integrate the two demands and adjust from the middle-ground. Anchors can exert an antagonistic effect, they pull the decision-maker in different directions rather than cancel each other out.

The reasonableness of the anchor influences the anchor effect. Implausibly high anchors, for example, have been found to have a proportionally smaller effect on judgements than anchors reasonably close to the expected value (Chapman & Johnson, 1994). Bahník and Strack (2016) suggest that the proximity between the anchoring value and real value may influence the occurrence of the anchoring effect.

Other empirical literature suggests that judges in small claims courts are more likely to split the difference between the litigants’ number than those in higher courts (Finkle & Cohen 1993).

[To be fleshed out]

The added value of the empirical study in this paper is the degree to which I can assess the strength of the underlying case. I have highly detailed information about the facts of the plaintiff’s case that indicate what the damages award should be. That is, I can provide a somewhat objective measurement of what a judge would award *if* the potential anchors were not present. Further, I can also generate a prediction of how other plaintiffs and defendants have behaved in other similar cases to generate an *expected* plaintiff claim and an *expected* defendant offer.

3. LEGAL SETTING AND HYPOTHESES

The legal background

I create a novel dataset that describes Canadian employment cases. The legal issue in each of my cases is the quantum of damages to be awarded in lieu of reasonable notice. Under Canadian law, workers who are dismissed without cause are entitled to a reasonable notice period. Statutory schemes in each province and territory provide a minimum number of weeks of notice that

workers are entitled to, based primarily on their length of service. But common law obligations of employers are typically greater than these minimum standards. The common law stipulates that the notice given must be reasonable.

Courts have held that “there can be no catalogue laid down as to what is reasonable notice in particular cases.” (*Bardal v. Globe & Mail*, (1960) 2 D.L.R. 140 (Ont. H.C.), per McRuer CJHC at 145.) Instead, what is reasonable will be determined with reference to: (1) the character of the work; (2) the length of service (tenure); (3) the age of the worker; and (4) the availability of similar work, with reference to the worker’s education and experience.

These factors are referred to in the case law as the *Bardal* factors. There are additional, non-*Bardal* factors that judges have been taken into account when assessing what is reasonable. For example, judges have looked to whether the worker was induced to leave their previous employer and subsequently dismissed after a short period of time.

Ultimately, what is at stake in these cases is a dollar amount. But, in case that involves assessment of reasonable notice periods, the litigants and the judges focus on the time period, rather than the dollar amount. That is, the issues are typically framed in terms of weeks or months (e.g., “the plaintiff is awarded 10 months’ notice.”) For consistency, I will only refer to the number of months’ notice, rather than the dollar amount of damages sought.

An example may be helpful. Take the case of *Barton v. Rona Ontario Inc.*, 2012 ONSC 3809. In this case, 64-year-old Kerry Barton was deemed to have been wrongfully dismissed by Rona, a large retail store. Barton had worked for Rona for three years and eight months, managing their store in Barrie, Ontario. Rona put forward two arguments. They first argued they had just cause to dismiss Barton. Alternatively, if they *were* found to have wrongfully dismissed Barton, they argued that the appropriate notice period should be no more than 4 months. Barton argued that he was wrongfully dismissed and suggested that 12 months was the appropriate notice period given his age, difficulty in obtaining similar work, and other factors. Each of the parties presented the court with a series of precedents that supported their assessment.

Hypothesis 1: Judges ignore the litigants’ numbers

The first hypothesis is that judges ignore the numbers put forward by the litigants. Under this theory, judges do not factor arguments made by the parties into their decision making. Essentially, the assessments provided are too noisy to provide any genuine of signal of what damages award is

reasonable. Instead, the judge bases her assessment on what she independently deems appropriate in the circumstances and based on precedents with similar fact patterns. This approach takes a cynical view of the arguments presented by the parties. Suppose -- purely hypothetically -- that in *Barton v. Rona*, the judge weighed up all the arguments of the two parties, ignored the numbers put forward by the two parties, and independently assessed the award to be 10 months' notice.

Hypothesis 2: Judges "split the difference" between the litigants' numbers

In *Barton v. Rona*, if the judge were splitting the difference, the reasonable notice period awarded to the plaintiff would be 8 months (the mid-point between 4 and 12.) Under this theory, the judge offers a compromise. Both parties gain and lose something. Neither party walks away from the courtroom knowing that the other side got exactly what they asked for. An argument in favour of this methodology is that the judge suspects that both sides are bringing offers and evidence that are largely favourable to them. The judge supposes that the appropriate level of damages lies somewhere in the middle of these two views. The split-the-difference approach can wash out the biases of the parties.

The problem, of course, is that if parties know that judges split the difference, then they have greater incentive to present offers and evidence that are even more biased. The split-the-difference approach does little to incentivize parties to present offers that are reasonable. As such, if the judiciary takes a split-the-difference approach, it may discourage parties from presenting reasonable claims and offers that would likely reduce the area of disagreement before coming to court, which in turn has an adverse impact upon the likelihood of settlement.

Hypothesis 3: Judges accept the more reasonable - or fairer - of the two parties' assessments of damages

A third hypothesis is that judges accept the more reasonable assessment by the two parties. Let's say that judges select one of the two numbers put forward by the litigants. Under this theory, the judge commits to picking the more reasonable of the two assessments. In our hypothetical situation where the judge in *Barton v. Rona* independently assessed the appropriate award to be 10 months' notice, if the judge were acting in accordance with this hypothesis, they would most likely award the 12 months asked for by the plaintiff. This award is presumably more reasonable than the 4 months put forward by the defendant. This strategy does not require the judge to find a suitable

compromise; indeed, committing to this strategy would forbid such compromise. And this strategy does not require the judge to select an outcome that is the “most reasonable”; rather, they only need to choose the “more reasonable” of the two options provided. If parties know that judges award damages in this way, then the parties are incentivized to reduce the area of disagreement themselves, since both parties wish to be seen as more reasonable than the other.

* * *

These three hypotheses are crude characterizations of how judges may determine damages awards. There are, clearly, other hypotheses that one could develop and potentially test. There are nuances to each of the three hypotheses – and clearly there are different definitions of what ‘reasonable’ or ‘fair’ means in this context. But each of these hypotheses are, in their own way, testable to some extent. In the next section, I describe the data in more detail and discuss the empirical strategy.

4. DATA AND EMPIRICAL STRATEGY

The dataset

I have 1,757 published cases from 1997 to 2019 in Canada that involve judgments about reasonable notice awards. For each case, a research team created a highly detailed matrix describing variables of interest for each of the cases.

First, I track information about the length of the notice period that the plaintiff employee asked for and what the defendant employer offered. I look for language in the opinion that expressly discusses what the parties think is reasonable. Such information is not present in every case. There are 601 cases where judge expressly notes *both* what the defendant employer thinks is reasonable *and* what the plaintiff employee thinks is reasonable.¹ It is these 601 cases that are analyzed in detail in this paper.² But the complementary cases

¹ There are actually 618 cases where this information is present. But in 17 of these cases, the plaintiff employee and the defendant employer essentially agree on what the reasonable period should be, and the contentious legal issue is over whether the reasonable notice period should be given at all. Because there is no dispute over the numbers, I remove these 17 cases from the analysis.

² There is additional information that one could incorporate into the analysis. First, in 562 cases, I have information about the initial notice package that was given to the employee by the employer. Next, there are also more cases where the judge notes the numbers put forward by

where I do not have full information about the litigants' numbers will also be useful for assessing the "reasonableness" of the awards.

The data are complicated somewhat by the fact that litigants sometimes put forward a *range* of reasonable outcomes, rather than simply one number. Plaintiff employees, for example, might say that a range of 20 to 24 months is reasonable. In 160 of our 601 cases, the plaintiff employee puts forward a range. Similarly, defendant employers commonly put forward a range. For example, they might suggest that a notice period in the range of 8 to 10 months is more reasonable. This occurs in 282 of the 601 cases. I track both the minimum and maximum of these ranges.

Second, the outcome of each case is tracked. **Figure 2** below shows the frequency of the outcomes in the 601 cases. As **Figure 2** illustrates, judges tend to gravitate towards awards framed in terms of months. Awards longer than 24 months are rare.

one party or the other, but not both. For the purposes of this paper, I have not incorporated these data because I am testing the comparison of the numbers put forward by both parties. Further, there is additional information one can glean from the discussion of the precedents that each party cites in their briefs. These data will be explored in future drafts.

In the example case, *Barton v. Rona*, the litigants put forward various numbers. The defendant employer Rona did not offer a notice period initially. This is because the employee was dismissed without cause. At the litigation stage, Rona submitted that, if they lose on liability, a notice period of 4 months would be an appropriate damages award. They presented the court with a series of precedents they deemed to be similar, all with outcomes in the range of 3 to 6 months. The plaintiff, Kerry Barton, sought 12 months' notice. He presented the court with three "relatively current" precedents, each with outcomes in range of 9 to 12 months. In my dataset, therefore, the *Barton* case is simply coded as an employer offer of 4 months and an employee claim of 12 months.

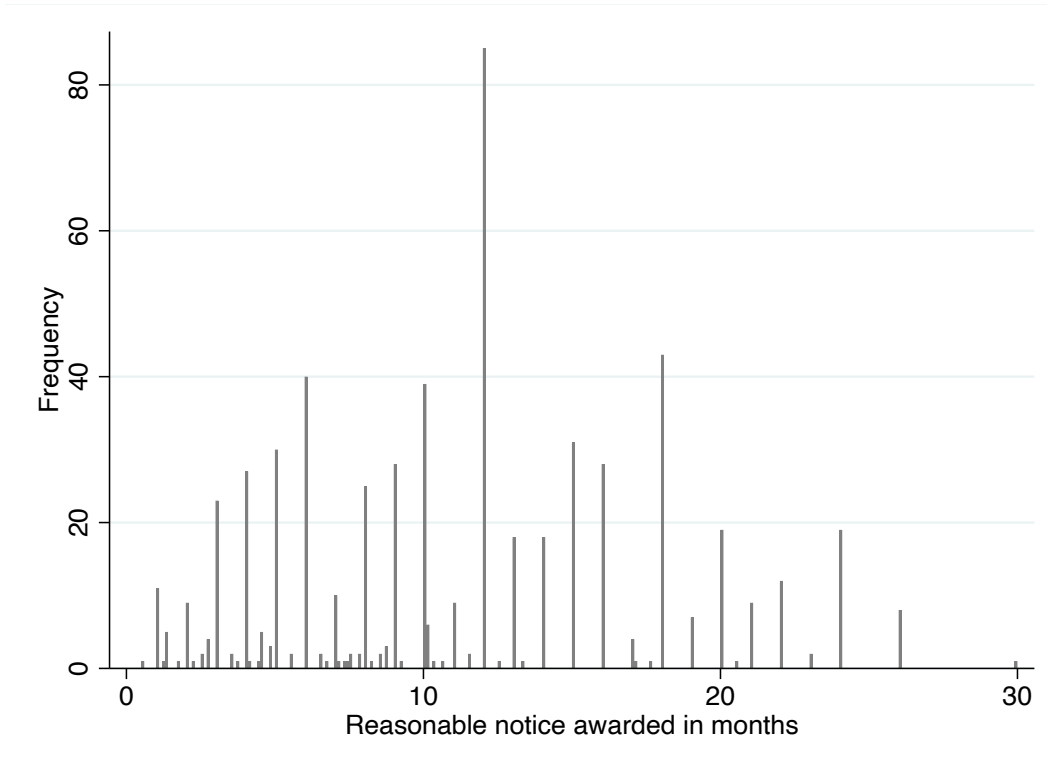


Figure 2: Damages, measured in terms of number of months’ notice, awarded by courts in 601 reasonable notice cases across Canada, 1997-2019.

Third, highly detailed information describing the facts of each case are collected. These data capture facts that are relevant to the interpretation of what is reasonable. The research team collected information for over 50 relevant variables pertaining to the *Bardal* factors as well as other relevant non-*Bardal* factors that judges consider. These data include information about the type of job, the level of seniority and responsibility of the worker, the compensation, the age of the employee, the length of service, the level of education and experience of the employee, factors relating the hiring and firing of the employee, and many other factors relating to job performance and the difficulty of finding future similar employment. Together, these data create a rich description of the relevant facts for every case that enable me to assess the strength of a plaintiff’s case.

The 601 cases are decided by 413 different judges. Most judges – indeed, the vast majority -- only decide one case in the dataset.³

³ There is one judge that decided 9 cases and three different judges that decided 6 cases.

Normalizing outcome data to “bargaining space”

For ease of comparing the outcomes of cases, I normalize the defendant employer’s low number, the plaintiff employee’s high number, and the ultimate judge’s award. I normalize the defendant employer’s low number to equal 0 and the plaintiff employee’s high number to equal 1. The area of disagreement in each of the 601 cases is thus normalized to a mass of 1.

The normalized outcome is measured as follows:

$$\text{normalized outcome} = \frac{\text{outcome} - \text{defendant offer}}{\text{plaintiff ask} - \text{defendant offer}}$$

If the judge’s award of damages is more favourable to the employer, then the normalized outcome will be a number close to 0; the more favourable to the employee, then the higher the normalized outcome will be, closer to 1. For example, in *Barton v. Rona*, the employer offered 4 months and the employee asked for 12 months. The court in this case ultimately awarded 10 months’ notice. In the dataset, the normalized outcome of *Barton v. Rona* is therefore coded as 0.75. The outcome is more favourable to the dismissed worker.

This approach is complicated by the fact that litigants sometimes put forward a range of numbers, rather than a single number. Suppose – purely hypothetically – that, in *Barton v. Rona*, the plaintiff employee asked for an award in the range of 12 to 18 months and the defendant put forward a range of 2 to 4 months.

To deal with this wrinkle, I create two different normalized measures for each case:

- (1) *largest area of disagreement*, using the lowest number put forward by the defendant employer and the highest number put forward by the plaintiff employee.
- (2) *smallest range of disagreement*, using the highest number put forward by the defendant employer and the lowest number put forward by the plaintiff employee.⁴

In this hypothetical example, the largest possible area of disagreement would be 2 months to 18 months. Here, the outcome (10 months) would be normalized in bargaining space to 0.5. The smallest possible area of disagreement would 4

⁴ I could, of course, also include the mid-points of these ranges. I do not include them here since I am looking to whether judges anchor to the actual numbers provided by the judges.

months to 12 months. The outcome (10 months) would be normalized in bargaining space to 0.75.

Generating a prediction of what is “reasonable” or “fair” based on the outcomes of precedents

I use a simple linear prediction model to predict the outcome of the cases based only on the facts of the case. I use data from the 1,156 cases where the judge does not expressly note the numbers put forward by both parties. For each of these cases, I have detailed information about the case, including the age of the employee, how long the employee worked at the firm, the type of job they performed, the availability of other work, the type of industry, compensation, the employee’s performance, along with many other salient facts. These variables are part of the legal test for determining what notice period is reasonable. I also include other variables such as the province in which the case was heard.

The statistical model describes and explain basic patterns in the case law. It allows me to determine the strength of the plaintiff’s case in the 601 cases. Essentially, the facts of each case are compared to the facts of every other case in the dataset, and – knowing the outcome of every other case – I can make a somewhat crude prediction of what the most likely outcome will be in the 601 out-of-sample cases if the litigant numbers were not presented. These predictions are independent of the numbers put forward by the two parties.

As noted above, the numbers put forward by the litigants are not factored into the prediction. In the example case of *Barton v. Rona*, the predictive linear model compared the facts of the case to the 1,156 cases used to predict the outcome. The model predicted an outcome of 6.4 months. This prediction did not consider the employer’s offer of 4 months or the employee’s claim of 12 months. The normalized prediction here was 0.3. This means that, according to this measure, the outcome was predicted to be closer to the numbers put forward by the defendant than to the numbers put forward by the plaintiff.

Generating predictions of what is “reasonable” or “fair” numbers by reference to the numbers of other litigants

An alternative way of exploring whether a litigants’ number is reasonable or fair is to compare it to the behavior of other litigants.

I regress the plaintiffs’ claims on the facts of cases. I then use a linear prediction model to generate an *expected* number that each plaintiff would put forward,

given the facts of their case. If the actual number put forward by the plaintiff exceeds this prediction, then they are acting more aggressively than expected plaintiff. If the actual number put forward by the plaintiff is lower than the prediction, they are acting more passively than other plaintiffs.

A similar exercise can be conducted with numbers put forward by the defendants. I regress the defendants' offers on the facts of the cases. I generate an *expected* number that a defendant would put forward. If the defendant in a particular case puts forward a number that exceeds this estimate, they are acting generously. If they put forward a lower number, they are acting more aggressively than expected.

5. PRELIMINARY RESULTS

5.1 The relationship between outcomes and litigants' numbers

Consider the outcomes of the cases relative to the numbers put forward by the litigants using the *largest area of disagreement*. To reiterate, the largest area of disagreement takes into the lowest number put forward by the defendant employer (normalized to 0) and the highest number put forward by the plaintiff employee (normalized to 1.) The relationship between outcomes and litigants' numbers is shown in **Figure 3**.

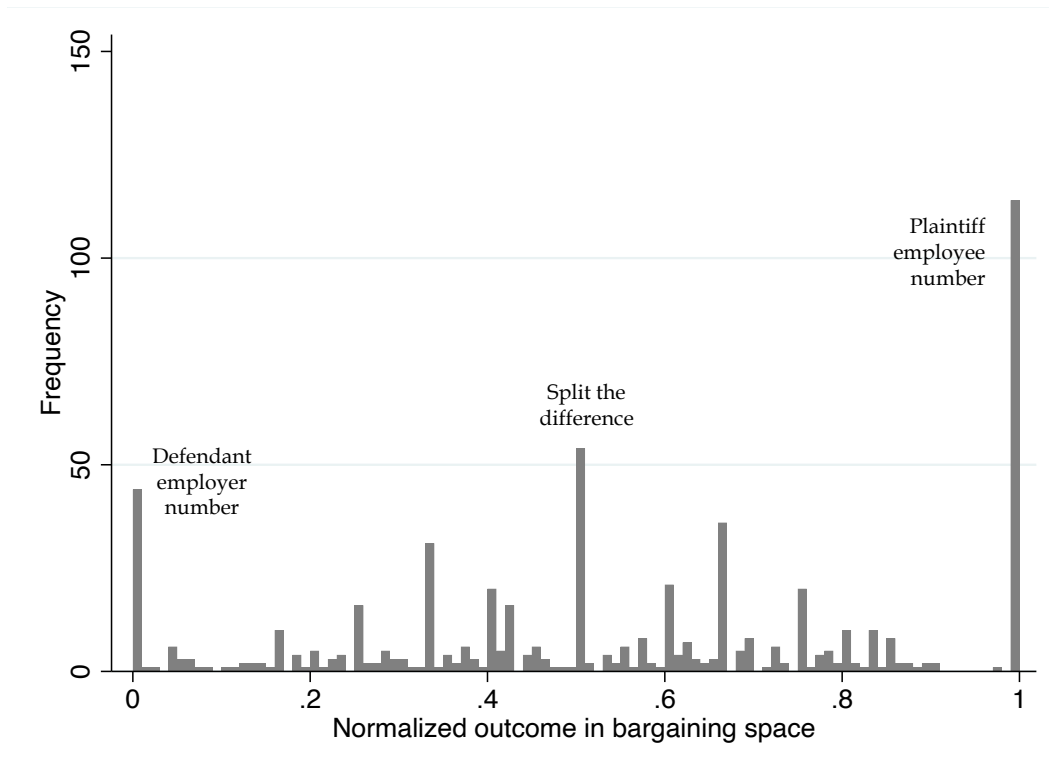


Figure 3: Normalized outcomes locating the reasonable notice period awarded within the *largest area of disagreement* between 0 (the lowest number put forward by the defendant) and 1 (the highest number put forward by the plaintiff.)

First, note that the modal outcome is 1. This is the outcome in 114 of our 601 cases (19.0%.) This means that the judge awards the highest number that the plaintiff asks for in about one in every five cases.

The second most common number in this spectrum is 0.5. This is the outcome in 54 of the cases (9.0%.) This means that judges select an outcome that splits the difference between the highest number asked for by the plaintiff and the lowest number offered by the defendant.

The third most common number is 0. This is the outcome in 43 cases (7.2%.) The judge in these cases awards the lowest number put forward by the defendant employer.

This picture changes somewhat when I normalize the bargaining space to the smallest area of disagreement (**Figure 4**). The smallest area of disagreements considers the highest number put forward by the defendant employer

(normalized to 0) and the lowest number put forward by the plaintiff employee (normalized to 1.)⁵

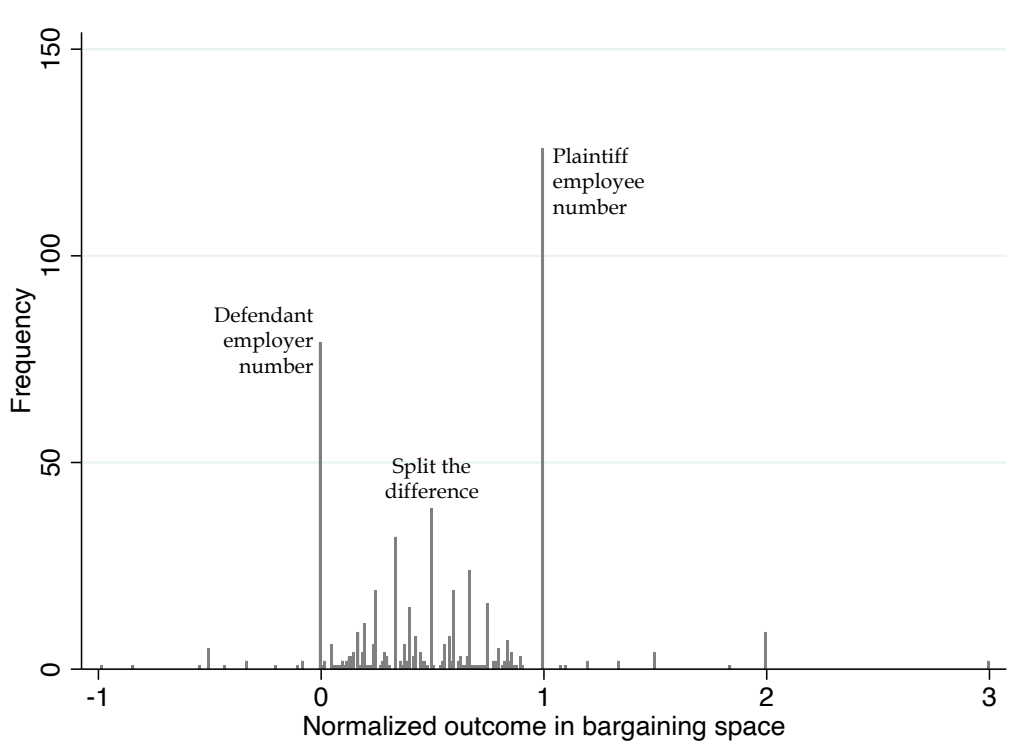


Figure 4: Normalized outcomes locating the reasonable notice period awarded within the *smallest area of disagreement* between 0 (the highest number put forward by the defendant) and 1 (the lowest number put forward by the plaintiff.)

There are cases where the judge awards notice periods that are longer than the minimum amount that the plaintiff put forward. For example, in *Carter v. Packall Packaging Inc.*, 2004 CarswellOnt 384 (ONSCJ), the plaintiff employee asked for 10 to 12 months, the defendant employer put forward a range of 7 to 9 months, and the court awarded 12 months. Measuring the area of disagreement as 9 to 10 months, a judicial award of 12 months puts the outcome outside of the area of disagreement. (This case is the outlier in **Figure 4**, with a normalized outcome of 3.)

⁵ There are 20 examples where the highest number in the range put forward by the defendant employer is the same as the lowest number put forward by the plaintiff. For example, the plaintiff in *Smith v. Ramara (Township)*, 2008 CarswellOnt 1909 (ONSCJ), the plaintiff sought a range of 12 to 13 months, while the defendant put forward a range of 10 to 12 months as being reasonable. Unsurprisingly, in this case, the judge decided that 12 months was reasonable. These 20 cases do not appear in the analysis when focusing on the smallest area of disagreement, because – by this metric – there is no area of disagreement.

There are also cases where the judge awards a notice period that is shorter than the highest number put forward by the defendant employer. For example, in *Ducsharm v. Quintette Operating Corp.*, 1998 CarswellBC 2846 (BCSC), the plaintiff asked for 12 months' notice. The defendant employer put forward a range 8 to 10 months. The judge awarded 9 months' reasonable notice. Since this number is lower than the highest number put forward by the defendants, the normalized outcome is less than 0.

Nonetheless, once again the modal outcome is 1. The most frequent outcome is that the judge awards the plaintiff the notice period that they were asking for (126 cases.) The second most common normalized outcome is 0, where the notice period is the one sought by the defendant employer (78 cases). The third most common is where the difference is split, with a normalized outcome of 0.5 (39 cases.)

Figure 1 (the figure presented in the Introduction) is somewhat of an amalgam of **Figure 3** and **Figure 4**. In **Figure 1**, the smallest area of disagreement is used. But if the judge awards a notice period that is lower than the defendant's largest offer, a score of 0 is assigned. If the judge awards notice period greater than the plaintiff's lowest number, a score of 1 is assigned. An outcome that splits the difference is recorded if the outcome is the midpoint of either the smallest area of disagreement or the largest area of disagreement.

* * *

In short, the judge awards damages that fall within the plaintiffs' range of suggested outcomes in 157 of the 601 cases. The judge awards damages within the defendants' range of suggested outcomes in 94 of the 601 cases. That is, in 41.76% of cases, the judge uses the numbers put forward by one of the two parties.

Further, in 78 cases (13.0% of cases), the outcome can be characterized as one in which the judge splits the difference between what the numbers put forward by the defendant and those put forward by the plaintiff (whether splitting the largest area of disagreement or the smallest.)

6. TESTING THE HYPOTHESES

6.1 Testing hypothesis 1: Do judges act independently of the numbers put forward by the parties?

The graphical evidence presented in Section 5 provides strong suggestive evidence that outcomes disproportionately reflect either of the parties' assessments – or they reflect the midpoint of these numbers. This provides *prima facie* evidence *against* the theory that judges simply ignore the litigants' numbers. If judges truly were ignoring the litigant numbers, one would not expect the actual outcomes to reflect them.

These preliminary results are, however, only suggestive evidence for this proposition. It may be, for example, that the judge ignores the numbers put forward by the parties and simply awards a notice period that she considers to be most reasonable, or one that is in line with the precedents. It may be that this number perfectly aligns with the numbers put forward by one of the parties.

There are additional tests one might employ to determine that the judges do indeed factor in the litigants' numbers. Let's turn to the linear prediction model. This model provides a prediction of the reasonable notice period that would be awarded if the judge only took into account the facts of the case and acting consistently with other cases.

First, consider how the litigant numbers provide an enhanced explanation of the ultimate outcome compared to the facts and the case law. Consider a linear regression of the outcome based only on the facts in the 601 cases. A regression that does not include the litigants' numbers has an *R-squared* of **0.812**. The variation in the facts of the cases can explain 81.2% of the variation in the outcome.

Let's now add in information about the litigant numbers to the prediction. If judges acted independently of the litigant numbers, the *R-squared* should not increase by much. For example, adding a constant independent variable (one that does not correlate with the outcome) to the linear prediction model would not increase the *R-squared* at all. Further, if the variation in the litigant numbers were largely explained by the facts, then adding in the litigant numbers should do little to increase the *R-squared*. If, on the other hand, judges typically anchor to either the plaintiffs' numbers or the defendants' numbers, then the *R-squared* would increase markedly. That is, more of the variation in the outcome could be explained by variation in the litigant numbers.

When I include the numbers put forward by the litigants into the linear prediction model, the *R-squared* increases to **0.867**. That is, 86.7% of the variation in the outcome can be explained by the facts and the litigant numbers. About 29% of the unexplained variation in reasonable notice awards when we only consider the facts can now be explained by looking at the variation in the numbers put forward by the litigants. This is an appreciable increase that

suggests that judges are using the numbers put forward by litigants as an anchor.

Another way to explore the added value of knowing the litigant number is to see how the accuracy of the prediction model improves when information about litigant numbers is included. Using just the facts of the case, the linear prediction model has an average prediction error of 2.01 months. When I include the numbers put forward by the litigant, this prediction errors falls to 1.67 months (a 17% improvement in accuracy.)

Finally, consider the distribution of predicted outcomes, normalized by the litigant numbers. If the predicted outcomes tended to coalesce around the 0 or 1 in the way that the actual data do (see Figure 2), then one may not be able to separate a theory of judges anchoring to parties' assessments from a theory of judges simply following precedent. The distribution of predicted outcomes, normalized by the litigant numbers, is presented in **Figure 5**. This figure illustrates that the predicted outcomes do not bunch at 0 and 1 in the same way as actual outcomes.

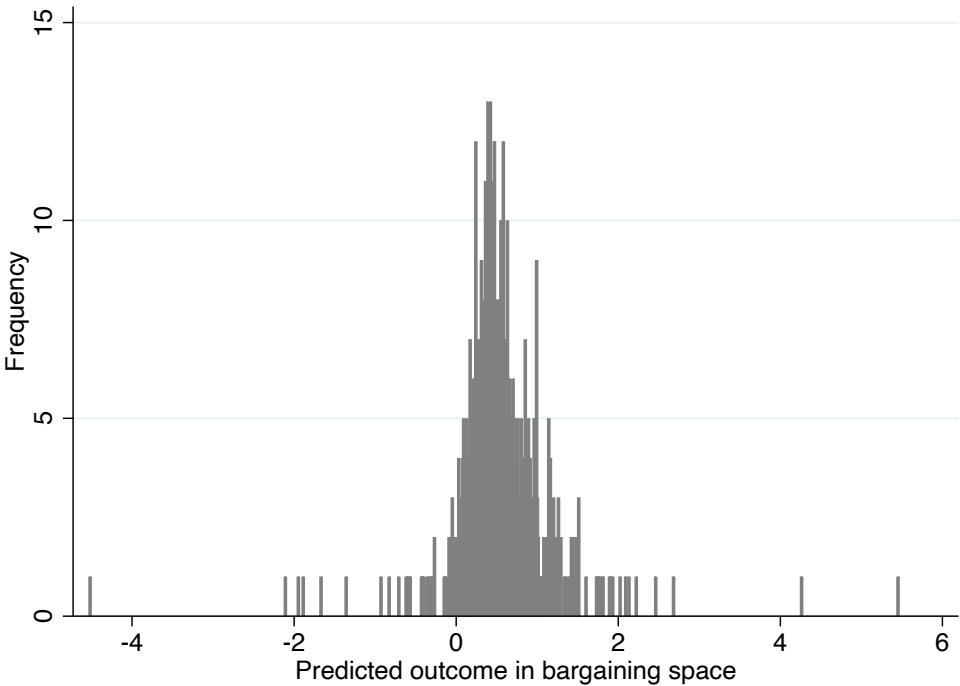


Figure 5: Distribution of predicted outcomes. The predicted outcomes have been normalized to fit within the area of disagreement where 0 is the lowest number put forward by the defendant employer and 1 is the highest number put forward by the plaintiff employee.

* * *

Taken together, this evidence suggests that judges are not ignoring the litigant numbers. Judges do not act independently of the numbers put forward by the plaintiff and defendant. If judges did ignore these numbers, we would not expect the inclusion of such numbers to increase the explanatory power or the predictability of the reasonable notice awards.

6.2 Testing hypothesis 2: Do judges “split the difference”?

Figure 3 and **Figure 4** offer evidence that judges split the difference in only a small subset of cases. As noted above, the outcome can be characterized as one in which the judge splits the difference between what the numbers put forward by the defendant and those put forward by the plaintiff (whether splitting the largest area of disagreement or the smallest) in 78 of the 601 cases (13.0%).

While not overwhelming in terms of numbers, a normalized outcome of 0.5 – one that splits the difference between the plaintiff’s high number and the defendant’s low number – is the third most frequent outcome overall.

But, even if the outcomes *do* coalesce around the midpoint, it may be that the predicted outcomes *also* coalesce around the midpoint. If that were case, it would be difficult to reject the theory that judges are simply acting in a manner that is consistent with following precedent.

To test this, I look at those cases where the predicted outcome (using facts of the cases only) is in the neighbourhood of the midpoint. There are 223 cases where the predicted outcome is in the neighbourhood of the midpoint. These are cases where the normalized predicted outcomes fall between 0.33 and 0.67.

Figure 6 shows the distribution of normalized outcomes when the predicted outcome is in the neighbourhood of the midpoint (left panel, 223 cases) and when the predicted outcome is outside of the neighbourhood (right panel, 378 cases). The proportion of cases where the judge “splits the difference” is marginally higher when the case is predicted to fall close to the midpoint (left panel). But the difference is not statistically significant.




Figure 6: Distribution of outcomes in bargaining space. Panel A (left) shows the distribution of outcomes in the 223 cases where the predicted outcome fell in the neighbourhood of the midpoint. Panel B (right) shows the distribution of outcomes in the 378 cases where the predicted outcome was outside of the neighbourhood of the midpoint. The proportion of outcomes where the award “splits the difference” is moderately higher in Panel A.

Judges split the difference in 13% of cases. The evidence that judges are simply following precedent in these cases is not overwhelming. It provides further evidence that judges use the litigants’ numbers as an anchor.

6.3 Testing hypothesis 3: Do judges gravitate toward to the “most reasonable” number presented by a litigant?

In about 42% of cases, the judge either awards the low number put forward by the defendant employer or the high number put forward by the plaintiff. But does that mean that judges gravitate to the number that is “most reasonable” or the “fairest” number?

6.3.1 Reasonableness as measured by reference to the predicted outcome

Let’s first define “most reasonable” in this context as the number that is most proximate to the predicted outcome (using the facts of the case only.) I ask: Is the predicted outcome closer to the defendant’s low number or closer to the plaintiff’s high number?

The evidence suggests that judges do gravitate to the more reasonable number. But the evidence is not overwhelming.

- Where the predicted outcome was closer to the plaintiff’s high number *and* the judge either selected one of the litigants’ numbers, the plaintiff’s high number was chosen in 75.7% of cases.

- Where the predicted outcome was closer to the defendant's low number *and* the judge selected either one of the litigants' numbers, the defendant's low number was chosen in just 61.8% of cases, slightly better than chance.

Judges are more likely to choose the more proximate litigant to the predicted outcome. The evidence suggests that judges do pick the "more reasonable" number more often than not if they decide to choose one of the litigants' numbers. But these numbers do not suggest uniform consistency in following this practice.

The evidence is somewhat stronger in the cases where linear prediction model generates a predicted outcome that falls outside of the [0,1] range. There are 72 cases in the dataset where the model predicts an outcome that is greater than the highest number put forward by the plaintiff. In these cases, the judge awarded the number put forward by the plaintiff in more than half of all these cases (38 cases, 52.8%).

There are 36 cases in the dataset where the model predicts an outcome that is smaller than the low number put forward by the defendant. The judge awarded the number put forward by the defendant in half of all these cases (18 cases, 50.0%). In none of these cases did the judge award the higher number put forward by the plaintiff.

6.3.2. Reasonableness as measured by estimated litigants' numbers

Now consider the alternative measure of the reasonableness of the litigants' numbers. I create an expected plaintiff claim by regressing plaintiff claims on the facts of each case. I then compare the estimate to the actual claim made. The difference between the two suggests how aggressive - or unreasonable - the plaintiff claim may be. If the actual claim exceeds the estimated claim, it suggests the plaintiff is more aggressive relative to her peer plaintiffs. I measure both absolute difference and proportional difference between the estimate and the actual claim.

There is a strong relationship between the aggressiveness of the plaintiff and the likelihood that the judge will anchor to the plaintiff's claim. **Table 1** summarize the data. For the 100 most aggressive plaintiffs, it is exceedingly rare for judges to anchor to the plaintiffs' number. Depending on how aggressiveness is measured, the plaintiff receives at least the minimum of their range in just 3 or 4 cases.

If plaintiffs are more passive compared to their peer plaintiffs, then it is more likely that the judge will anchor to their number. For the 100 most passive plaintiffs, judges anchor to their number in about half the cases (55 of the 100 most passive cases as measured by absolute difference; 48 of the 100 most passive where measured by proportional difference.)

Proportion of cases where judge anchors to plaintiff claim

	<u>Rank by absolute difference</u>	<u>Ranked by proportional difference</u>
100 most aggressive plaintiffs	0.03	0.04
200 plaintiffs ranked 101 to 300	0.22	0.22
201 plaintiffs ranked 301 to 501	0.35	0.38
100 most passive plaintiffs	0.55	0.48

Table 1: Proportion of cases where the judge anchors to the plaintiff claim, broken down by subsets based on the aggressiveness of the plaintiffs’ claim.

Turning to the defendant’s offers, the data tell a similar story. The lower the defendant’s offer, relative to what is expected, the less likely the judge is to anchor to their number. This is shown in **Table 2**. The relationship is strikingly similar to that of plaintiff offers. The judge anchors to the defendant’s offer in just 2 or 3 (depending on the measurement) of the 100 cases involving the most aggressive defendants. Judges are more likely to gravitate toward the defendant’s offer as the offer becomes relatively more generous. For the 100 most generous defendants, the judge anchors to their higher number in nearly half the cases.

Proportion of cases where judge anchors to defendant offer

	<u>Rank by absolute difference</u>	<u>Ranked by proportional difference</u>
100 most aggressive defendants	0.02	0.03
200 defendants ranked 101 to 300	0.13	0.12
201 defendants ranked 301 to 501	0.22	0.20
100 most generous defendants	0.42	0.47

Table 2: Proportion of cases where the judge anchors to the defendant offer, broken down by subsets based on the aggressiveness of the defendant’s offer.

The evidence suggests that judges are more likely to anchor to the number put forward by the litigant if it is more reasonable, measuring reasonableness in relation to the expected number. This result is intuitive.

Interestingly, though, there is no statistically significant relationship between the aggressiveness of a plaintiff's claim and the likelihood that the judge will anchor to the defendant's offer. Similarly, there is no relationship between the aggressiveness of a defendant's offer and the likelihood that the judge will anchor to the plaintiff's claim.

Thus, while there is strong evidence that the reasonableness of the number will affect the chances that the judge will anchor to that number, there is little evidence to suggest that the judge will simply gravitate to the opponent's number.

7. DISCUSSION / CONCLUSION

[To be completed]