



A simulation model of the NSW District Criminal Court with illustrative applications

Hamish Thorburn

Aim: To determine the effect of various possible reforms on the pending case backlog in the Sydney registry of the NSW District Criminal Court using a simulation model. Specifically to determine what (if any) interventions can reduce the case backlog to 430 cases or fewer by December 31 2019.

Method: A simulation model was developed using the R statistical software package, and was used to determine the likely effect of various initiatives in reducing trial court delay.

Results: If no preventative measures are taken, the backlog will rise from the July 2016 level of 702 cases to approximately 796 cases by December 2019. The only intervention which will reduce the case backlog to 430 cases or fewer by December 2019 is the addition of 5 or more judges (a 33% increase), which would reduce the backlog to (on average) 375 cases by December 2019. A 75% reduction in late guilty pleas would result in the backlog falling to 725 cases, any reduction in trial-day adjournments would have virtually no effect, and the elimination of the mid-year court vacation would result in the backlog falling to 648 cases. A combination of 2 additional judges (an 11% increase), a 50% reduction in late guilty pleas, a 33% reduction in adjournments and the elimination of the mid-year vacation periods would also reduce the backlog to (on average) 564 cases by December 2019.

Conclusion: The only intervention which could reduce the backlog to 430 cases or less by December 2019 is the addition of 5 or more judges to the Sydney registry of the NSW District Criminal Court. A combination of 2 additional judges (an 11% increase), a 50% reduction in late guilty pleas, a 33% reduction in adjournments and the elimination of the mid-year vacation would achieve a smaller but still significant reduction.

Keywords: District court backlog, trial delay, simulation model.

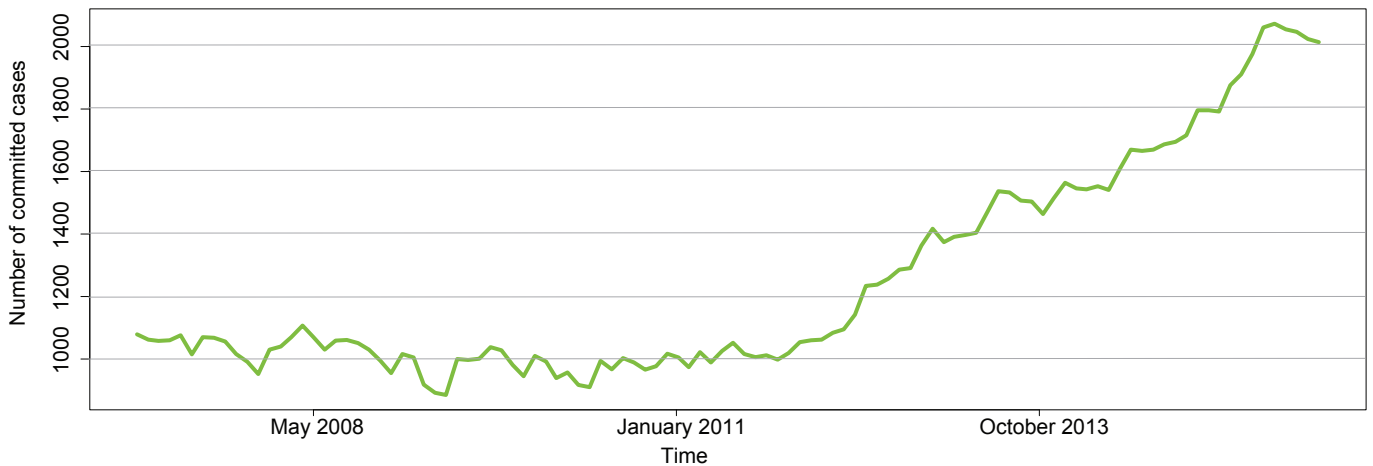
INTRODUCTION

The prison population of New South Wales (NSW) is currently at record levels. In June 2016, the population reached 12,550, a 12% increase over the population for June 2015 (NSW Bureau of Crime Statistics and Research, 2016). This rise has been attributed in part to a large growth in the number of prisoners on remand (that is, prisoners awaiting trial or sentence), which rose by 14.8% between July 2015 and June 2016 (compared to a 2.8% rise in the sentenced prisoner population over the same period – see [NSW Bureau of Crime Statistics and Research, 2016]).

One of the reasons behind this rise is an increase in court delay (that is, the number of days between committal to trial from the

Local Court and case finalisation in the District Court). District Court data shows that the median court delay in the NSW District Criminal Court (DCC) has risen from 223 days in 2011 to 348 days in 2016¹. This delay was examined more thoroughly by Weatherburn & Fitzgerald (2015). They identified a growth in the pending trial caseload in the NSW DCC, caused by an increase in the number of registrations (i.e. committals from the local court) relative to the number of finalisations and a slight increase in average trial duration. Furthermore, there has recently been a substantial growth in the number of cases awaiting trial in the DCC. Figure 1 shows the pending caseload of the NSW DCC each month between January 2007 and July 2016. The caseload has risen from roughly 1,000 cases in 2011 to 2,042 cases in July 2016. It should be noted that this does not mean that more

Figure 1. Number of committed cases awaiting trial in the NSW DCC vs Month, January 2007 to July 2016



than 2000 cases are waiting for trial. The majority of cases in the pending trial caseload will actually end in a plea of guilty rather than a trial. There is, nonetheless, a definite relationship between the size of the pending trial caseload and the time taken to finalise matters that do proceed to trial (Chilvers, 2001).

This is not the first time that NSW has experienced issues with court delay. In 1989, the NSW Bureau of Crime Statistics and Research (1989) reported that the delay between committal and finalisation in the District Courts ranged between 14 and 26.5 months for different matters and that, in the Higher Courts, the delay was due to an increase in committals from the Local Court throughout the 1980s. A number of recommendations were put forward to reduce the backlog (Coopers & Lybrand WD Scott, 1989), including the removal of committal proceedings, permitting audio/video recordings of confessional evidence, changes in court administration procedures, stricter requirements for case adjournments, the appointment of acting judges, a lengthening of court days to deal with short matters and a possible reduction in the January 'vacation' of the courts² (see Coopers & Lybrand WD Scott [1989] for a full summary of recommendations). In response to this, many of the proposed recommendations were adopted (including audio recordings of evidence and the appointment of additional judges, as well as other changes which were not suggested by the report) and court delay fell from a median of nearly 600 days in 1991 to less than 400 in 1996 (Weatherburn & Baker, 2000).

However, by 1999, median court delay rose again, reaching nearly 500 days. This rise was highlighted by Weatherburn & Baker (2000), who also examined possible reasons for the trend. They concluded that, rather than a lack of available court capacity, the growth in trial court delay was due to the high proportion of trials that failed to proceed on the first listed trial date (whether it be due to an adjournment or a late guilty plea)

and a policy of 'over-listing' trials (listing more trials than there is available court time to hear) as a strategy to combat this. The 'over-listing' policy, they argued, meant that both the Crown and the defence were often less prepared on the listed trial date (as they suspected they would not be reached), and hence, if called, would adjourn to a later date, further exacerbating the problem. The authors suggested that an increase in early guilty pleas and more certainty with regards to trials commencing on their designated start date were the best ways of trying to reduce court delay. At this point, median court delay began to drop again, falling to just 203 days in 2003 (Kuan, 2004). Since that time, court delay and the pending case backlog have risen once more – and in the case of the pending case backlog, to record levels.

Reducing the pending trial backlog is a priority for the NSW government. In December 2015, the Attorney General announced that the government would commit \$20 million in 2016 to fixing this problem (Needham, 2015). The funding was directed at creating extra sitting weeks in certain registries, appointing additional public defenders and temporary judges, and implementing measures designed to encourage early guilty pleas (where appropriate). In the 2016-2017 NSW state budget, a further funding package of \$39 million was set aside to help reduce the trial case backlog (NSW Government, 2016a).

Given the effort and resources currently being invested in reducing the backlog, it is appropriate to develop methods to assess the relative effectiveness of different delay-reduction options. One approach is to construct a simulation model of the NSW District Criminal Court and use the model to examine the likely effects of different policy scenarios. In this bulletin we present such a model and illustrate its use by exploring the effects of five different scenarios that could be employed to reduce court delay:

1. An increase to the number of judges available to preside over trials in the DCC.
2. A reduction in the number of late/trial-day guilty pleas³ by defendants in the DCC.
3. A reduction in the number of adjournments occurring on the morning of a trial in the DCC.
4. Eliminating the mid-year vacation period of the DCC.
5. A combination of scenarios (1) through (4).

In the next section we conduct a brief review of the use of simulation models in the criminal justice system. In the section that follows we explain the nature of the model and how it was constructed and validated. In the third section we present estimations of the effects of options (1) to (5) above on the trial case backlog of the NSW District Court. In the final section we discuss our findings and highlight some of the caveats surrounding our conclusions.

QUANTITATIVE MODELLING IN CRIMINAL JUSTICE RESEARCH

There are two general approaches to quantitative modelling of court systems. Firstly, there are queuing models, which use known results from queuing theory (a branch of mathematics) to determine analytical solutions to problems. These models take the form of a collection of mathematical equations describing either the number of customers (in our case, defendants) waiting to be served at any point, or the expected waiting time for each customer (see Greenberg, 1979). The second type of approach involves performing computer simulations of the process of interest, known as simulation models. Simulation modelling seeks to answer 'what if' questions, such as what would happen to court delay if the number of judges were increased or more people pleaded guilty. A forecast or baseline is first obtained (from a suitable model) predicting the likely course of (say) court delay if nothing changes. Simulation modelling is then used to see what would happen to court delay (relative to the baseline) if some factor influencing it were changed. Some models involve assumptions about the probability of each individual case/individual moving from one part of a system (their 'state') to another and the length of time spent in any one part or 'state'. These models are known as discrete event models. Of course, mathematical and computer simulation models are only as reliable as the assumptions on which they are based, but in a way that is one of their advantages. They allow the user to change the assumptions surrounding a process and see how much this alters the outcome. Their value lies less in giving a precise prediction of the impact of a policy on an outcome than in providing a 'confidence interval' around the effects.

A number of computer simulation models have been developed to assess the likely impact of changes in criminal justice

policies and practices. A few examples, presented here, serve to elucidate some of the applications of simulation modelling in the Australian criminal justice context. Crettenden, Packer and Macalpine (1993) developed a computer simulation model to look at case flows through the NSW DCC. The model allowed users to alter (among other things) the available court time, the number of registrations each month, the probability of a case being finalised at various points (to simulate late guilty pleas, adjournments, no bills, etc.), case priority rules and trial duration. The model was able to accurately replicate caseloads and waiting times for DCC trials in 1989, and hence could be used to examine the effects of increasing the court-time by various amounts. Worrall (1982) created a similar model for the Adelaide Magistrates Court.

Quantitative models have been used in other areas of the justice system as well. These include overall models of the criminal justice system (CJS) (showing, for example, how individuals move through the non-offending, remand, bail and prison populations), police models (often relating to duty/patrol scheduling), and corrections models (showing the flow of people through the corrections system) as well as many others (see Nagel 1977; Lind, Chilvers and Weatherburn 2001; Livingston, Stewart and Palk 2006). Lind, Chilvers & Weatherburn (2001), also, developed a stock-and-flow model of the criminal justice system to examine the effect of changes to various processes of this system on the local court bail, local court custody, district court bail, district court custody and prisoner populations.

METHOD

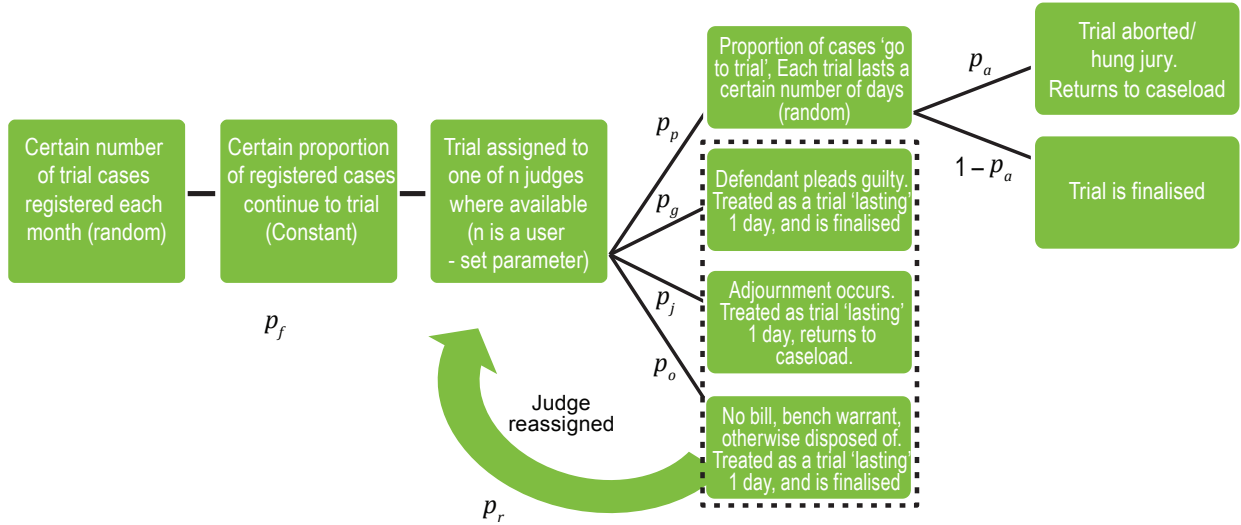
As the model developed by Crettenden, Packer and Macalpine (1993) is somewhat out of date, we present a new model of the DCC and describe the data and methods used to estimate its parameters. Initially, we attempted to model the entire state of NSW (as in, aggregating committals, finalisations and the backlog across all registries) from 2007 onwards. However, due to considerable variations in the backlogs, committal distributions and trial duration distributions across registries, this was impossible. Instead, we modelled only one registry - the registry of Sydney, which accounted for 33% of all finalised trials in 2015.

MODEL DESIGN

The basic model design is shown in figure 2.

The model works by sequentially stepping through weekdays over a pre-set time period. On each day (let t denote the current day), there are certain processes that take place. On the first day of each month, a certain number of trials to be committed from the Local Court are randomly drawn from a (rounded) normal distribution⁴ with mean A_m and variance σ_m^2 for month m . These trials have a probability of proceeding to the 'first day of the trial' with probability p_f and a probability of dropping out of the system

Figure 2. A flow diagram describing the NSW DCC model



of $1 - p_f$. The process then considers assigning the trials to the n available judges (n is a user-set parameter)⁵. Of the n judges, only some may be available to work on a trial. Others may be working on previous trials, on leave/non-sitting duties, or may be assigned to work on ‘short-matters’ (appeals and sentencing matters) each day.

Once the number of available judges (n') has been determined, trials are then assigned to these judges. Each trial then proceeds to one of four outcomes. The trial can proceed as normal, which occurs with probability p_p . In this case, a trial duration (d) is drawn from the empirical distribution of trial durations. That is, one of the trials from our sample dataset is randomly selected, and our new simulated trial is assumed to have the same duration as the trial selected from our sample. In this case, our sample was every trial heard in the NSW DCC between 2011 and 2016, which we obtained in a dataset from the NSW DCC. The trial is then said to be aborted (with probability p_a) and returned to the caseload, or finalised (with probability $1 - p_a$) and exits the system. If the trial does not proceed, then either the judge is reassigned to another trial (with the same possible outcomes) with probability p_r , or it is assumed that the vacancy cannot be filled and the day is wasted. This is effectively the same as a trial of one-day’s duration. If the judge is reassigned, then we assume no time is lost during reassignment. This is effectively the same as a trial of one-day’s duration. For simplicity, we assume that a judge can be reassigned at most once – if the new trial also does not proceed, then the judge is not reassigned to a new trial a second time. There are three different reasons that a trial may not proceed on its listed date. With probability p_g the defendant pleads guilty, and the case is disposed of as a ‘sentence matter’. With probability p_j an adjournment occurs, and the matter returns to the caseload. And with probability p_o the case is disposed of by some ‘other’ means⁵. The availability of the judges is then updated, t is

increased by one day, and the whole process repeats. The process terminates when t reaches a pre-specified date. The assumed values for the parameters are included in Appendix 1.

KEY MODEL ASSUMPTIONS

Some additional assumptions have been made in order to explore the effect of different policies, such as:

1. Whenever a judge is available to preside over a trial, a trial is ready to be heard. While both prosecutors and defendants need time to prepare their case, for the current levels of court delay, it is assumed that, at any given time when a judge becomes available, there is a least one trial case ready to proceed. This assumption becomes more tenuous at lower levels of the case backlog.
2. Apart from seasonal variation, there is no change in the average trial duration or average number of committals per month. We assumed no temporal change in trial duration due to insufficient data to use the empirical distribution for seasonal and temporal differences in trial duration. We assumed no temporal change in the monthly number of registrations using an ARIMA model – see Appendix 2 for details. This assumption may be open to question if the monthly number of trial registrations or some other factor (e.g. late changes of plea) results in longer trials.
3. The probability of cases reaching the morning of the trial (p_p), and the probabilities for various trial-day outcomes (p_p, p_g, p_a & p_o) remain constant over the entire simulation (aside from user-determined interventions).

PARAMETER ESTIMATION

Two datasets were obtained from the NSW District Court to estimate the parameters of the model. The first dataset contained the number of registered, finalised and pending trials

in each of the seven registries in the district court for each month between January 2007 and July 2016. This was used to estimate the frequency and distribution of cases being committed from the local court to the Sydney DCC (that is, how often cases are coming into the DCC) on a monthly basis. The second dataset contained information on the duration of each trial (in days) split by registry between January 2011 and June 2016. This was used to determine the distribution of trial durations in the Sydney DCC. Additionally, other data was sourced from the NSW District Court Annual Review (NSW Government, 2016b) and from members of the Office of the Director of Public Prosecutors (ODPP) (Lacey, 2016) in order to determine the transition probabilities (p_f, p_p, p_g, p_a & p_o) for the model. The NSW District Court Annual Review provided data on the 2015 trial listing outcomes (i.e. what happens to each trial on the day it is supposed to commence) and the ODPP data described how each trial was eventually finalised (regardless of how many times it was adjourned/ when the trial was finalised). By taking the probability that a case is adjourned on the first day of the trial (p_a) from the District Court Annual Review, and (for example) the probability that a case is eventually finalised by a guilty plea (which we will denote as p_{fg}), we can work out the probability that a case will be finalised by a guilty plea on the morning of a trial (p_g), which we find to be $p_g = p_{fg} (1 - p_a)$. Full derivation of this result (and similar results for other parameters) can be found in Appendix 3.

MODEL VALIDATION

An important step in developing any model is validating assumptions to ensure that the conclusions drawn from the modelling are realistic and to determine whether any particular aspect of the court process had been modelled inappropriately. The Sydney DCC model was validated in two ways. Firstly, the assumption of the number of committals being normally

distributed was assessed using a Shapiro-Wilk test for normality. Secondly, the case backlog was simulated 50 times and the current conditions of the Sydney DCC (current mean registrations, trial durations and numbers of judges), and a 95% prediction interval for the backlog at each month, were calculated. The observed case backlog was then checked to see if it was within this prediction interval.

VALIDITY OF THE REGISTRATION DISTRIBUTION

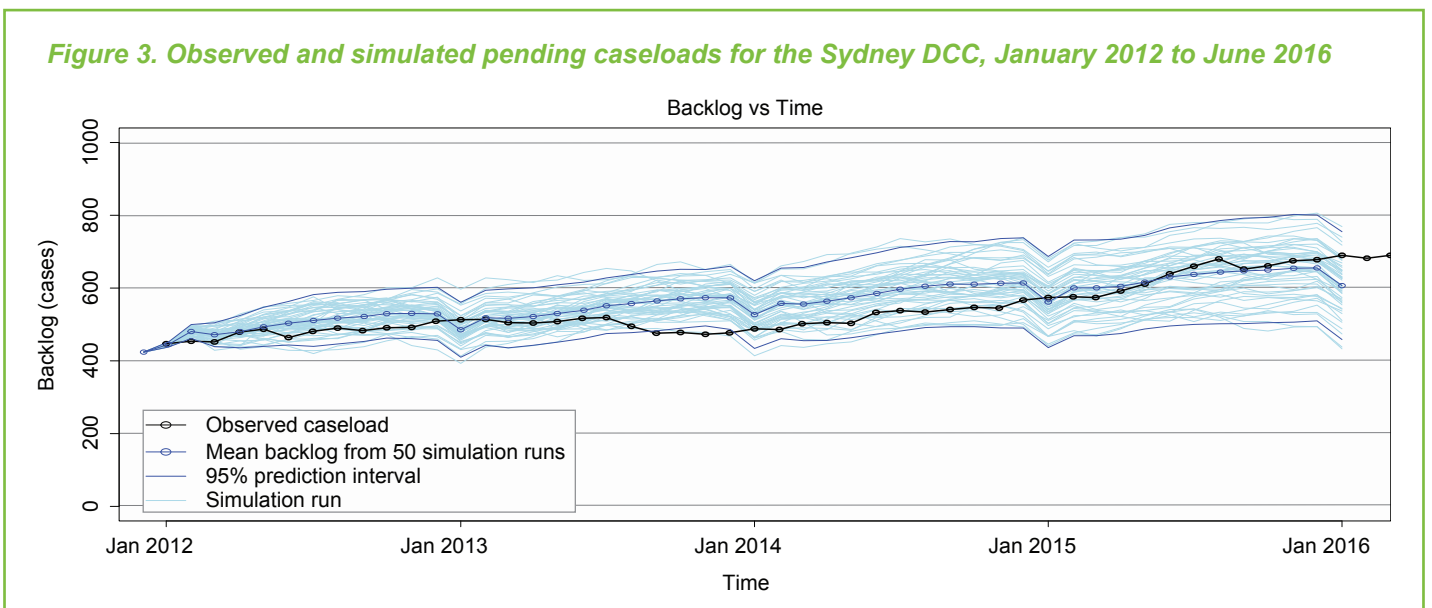
As noted earlier, the number of cases being committed to trial from the local court each month was assumed to follow a (rounded) normal distribution. We assumed that the registrations between the months of February through to November followed a common normal distribution (justification of this is included in Appendix 4). December and January were excluded as considerably fewer registrations were observed for these months. We showed that the normality assumption holds – see Appendix 4 for more details.

MATCHING THE OBSERVED AND SIMULATED PENDING CASELOADS

After verifying the distributional assumptions, we checked whether the pending caseload outputted by our model was similar to the observed pending caseload. We did this by performing 50 simulation runs of the case backlog with all parameters set to levels matching the current state of the Sydney DCC, simulated over the period of January 2012 to June 2016. A 95% prediction interval was found using these simulation runs, and the observed backlog was checked to see if it passed through this prediction interval, shown in Figure 3.

Figure 3 shows the number of cases in the backlog (y-axis) over time (x-axis) for 50 simulated runs of the model, as well as the observed Sydney backlog between January 2012 and

Figure 3. Observed and simulated pending caseloads for the Sydney DCC, January 2012 to June 2016



December 2015. Each light-blue line represents the backlog over time for one simulation run. The smooth blue lines show the 95% prediction interval for the backlog, and the blue line with circular markers shows the mean backlog. As can be seen in figure 3, that the observed backlog (black line) remains within the 95% prediction interval (smooth blue lines)⁶ and stays close to the mean backlog (blue line with circular markers). This simulation was run with n set at 18 judges (in reality, the Sydney registry has between 18-22 sitting each day) and a probability of reassignment (p_r) of .5. This shows that the simulated backlog closely matches the observed backlog, validating the model.

SIMULATION RESULTS

Note that all results shown here are based on the assumptions stated above, some of which could not be validated. What follows, therefore, should be taken as illustration of how the model can be used rather than a complete analysis of the relative merits of various options for reducing court delay.

Having validated the model to the best of our ability, we used it to simulate various interventions to determine the effect on the size of the case backlog. Given that the pending trial backlog in the Sydney District Court registry as of the end of July 2016 was 702 cases we ran our interventions from August 1 2016, aiming to reduce the pending caseload in Sydney from 702 cases to 430 cases (the average level of the pending caseload in Sydney in 2011).

First, we determined the baseline – that is, what happens to the backlog if no action is taken.

Figure 4 shows what happens to the backlog if no intervention is performed. The black dotted line shows the observed case backlog, the light blue lines each show the simulated case

backlog from one simulation run, the blue dotted line shows the mean backlog from the 20 simulation runs, the smooth dark blue lines show the 95% prediction interval for the backlog, the red vertical line shows the date in which any implemented intervention comes into effect, the horizontal dashed black line shows the target level for the backlog (430 cases) and the solid vertical black line shows the target date to reach the target caseload (December 2019). Figure 4 shows that the mean backlog continues to rise and reaches nearly 800 cases by 2020.

POSSIBLE INTERVENTION: INCREASING THE NUMBER OF JUDGES

An obvious intervention to consider is increasing the number of judges available to hear trials in the DCC. Note that when the number of judges is increased, it is assumed that there are available courtrooms, prosecutors and jury for additional trials to take place. The results of this intervention are shown in Figure 5.

Figure 5. Size of case backlog vs. time when introducing different numbers of judges for 20 simulation runs. The grey horizontal line shows the target backlog of 430 cases, and the black vertical line shows the target date of December 31 2019.

The graphs in Figure 5 show the projected size of the pending case backlog (y-axis) over time (x-axis) for different scenarios. In each panel, the solid black line with markers shows the observed case backlog in Sydney between January 2012 and July 2016. Each light blue line shows the average projected backlog from August 2016 to August 2021 for one simulation run. The average backlog of the simulation runs at each point is depicted by the blue line with circular markers, with 95% prediction intervals shown for this period by the smooth blue lines.

Figure 5a (which equates to the baseline shown in Figure 4) shows the expected size of the backlog of pending trials if no

Figure 4. Size of case backlog vs Time when no interventions are performed – the baseline case

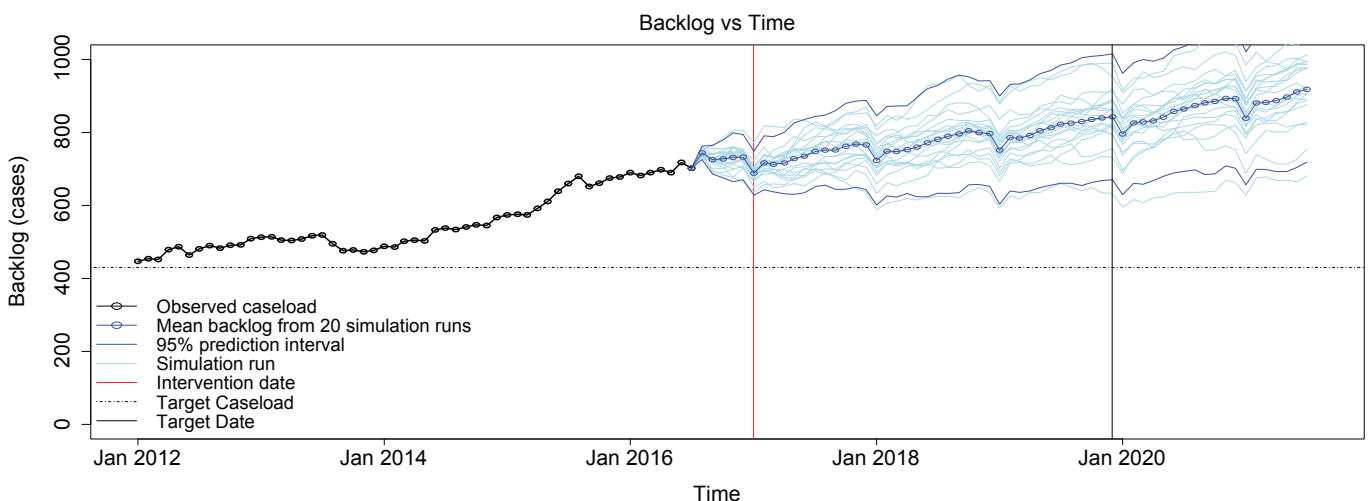
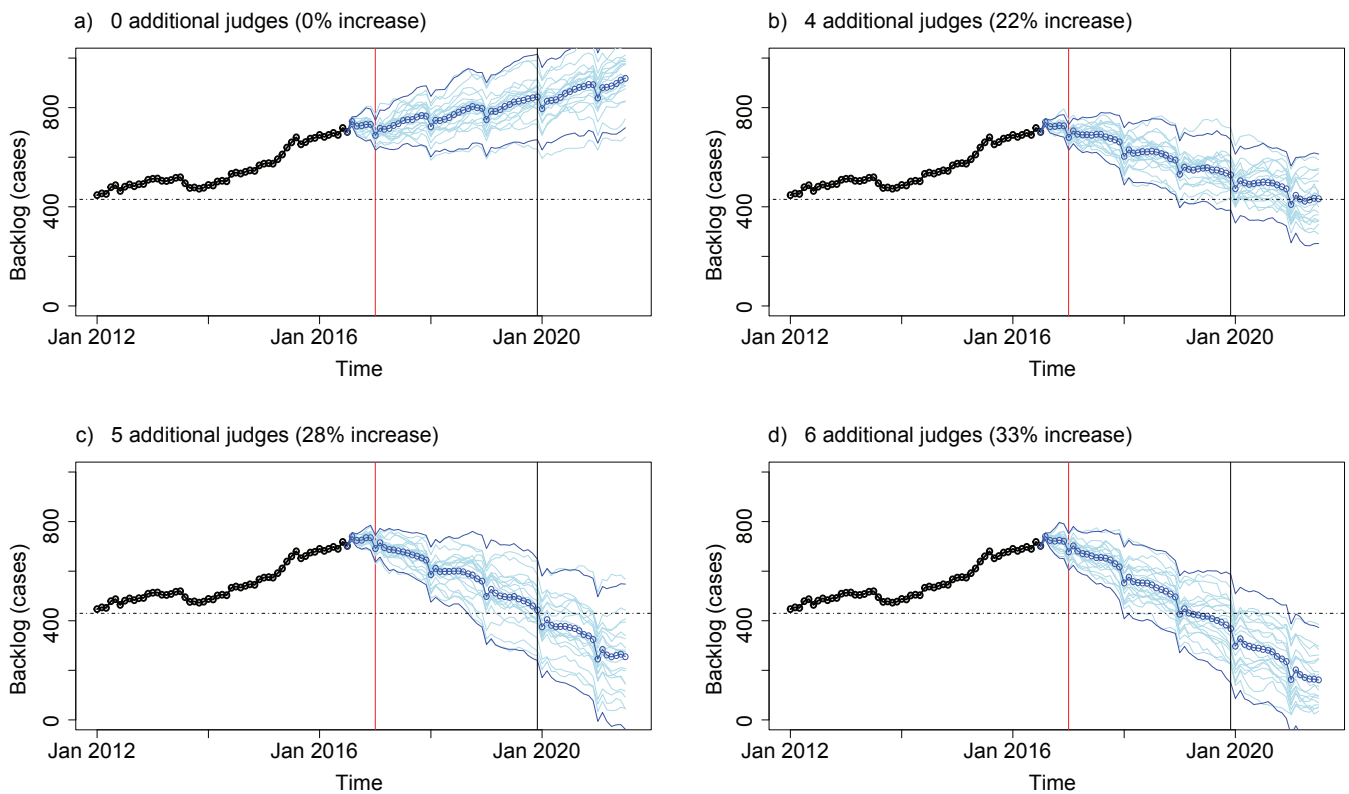


Figure 5. Size of case backlog vs. time when introducing different numbers of judges for 20 simulation runs



additional judges had been appointed to the District Court. Figures 5b to 5d explore the effect of adding four additional judges (Figure 5b), five additional judges (Figure 5c) and six additional judges (Figure 5d) in January 2017 (shown by the vertical red line). As can be seen, the more judges are added, the greater the effect on the projected backlog (as one would expect). The scenarios which results in the backlog falling to 430 cases (horizontal grey line) by the target date (black vertical line), however, is the addition of 5 or 6 judges. This results in an average case backlog by December 2019 of 375 cases for 5 judges and 297 cases for 6. The scenarios of 0 and 4 judges result in average backlogs in December 2019 of 796, and 473 cases respectively. It is worth noting that, in the scenario with 4 additional judges, while the average backlog doesn't reach 430 cases by the target date, 430 cases is contained within the 95% prediction interval at this point.

POSSIBLE INTEVENTION: REDUCTION IN TRIAL-DAY GUILTY PLEAS

Another strategy often put forward involves reducing late guilty pleas (i.e. those occurring on the morning of the trial and wasting valuable trial court time).

The reform whose effect we will examine here is one in which we assume $x\%$ of all guilty pleas which occur on the morning

of the trial and ALL guilty pleas which occur at arraignment (a court hearing between committal and trial) will now occur before committal to the District Court (i.e., they will plead guilty before entering the system). This is to replicate some proposed reforms as indicated by Hetheron (2016). In terms of the model, this involves reducing the probability that a defendant pleads guilty on the morning of the trial (p_g), as less cases are pleading guilty at this point. However, the defendants in these cases are not requesting adjournments or proceeding to trial instead – they are pleading guilty at committal, and hence not even entering the DCC system. This means that we also need to reduce the average number of cases being committed to the DCC (A_m). Furthermore, since the cases that would be pleading between committal and trial are now pleading at committal, we also need to adjust p_j . Details of how an $x\%$ reduction in late guilty pleas would change each of the aforementioned parameters are given in Appendix 3.

The results for various reductions in late guilty pleas are shown in Figure 6. The effects of a 25% (Figure 6b), 50% (Figure 6c) or 75% (Figure 6d) reduction in late guilty pleas compared to the baseline level (Figure 6a) are shown. It can be seen that even the most effective scenario (a 75% reduction in late guilty pleas (Figure 6d)), does not result in the backlog decreasing – the average backlog still rises to 725 cases in December 2019

Figure 6. Size of case backlog vs. time for different reductions in late guilty pleas for 20 simulation runs

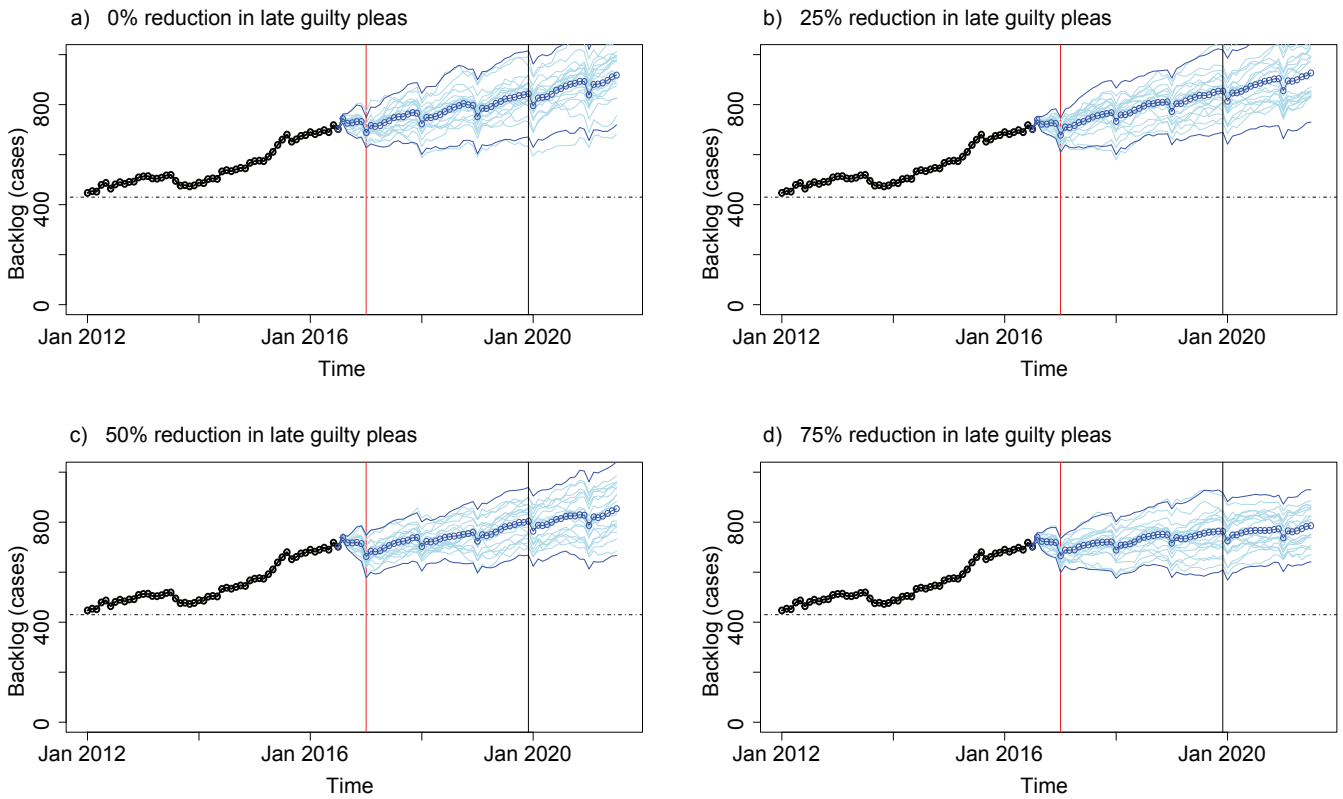
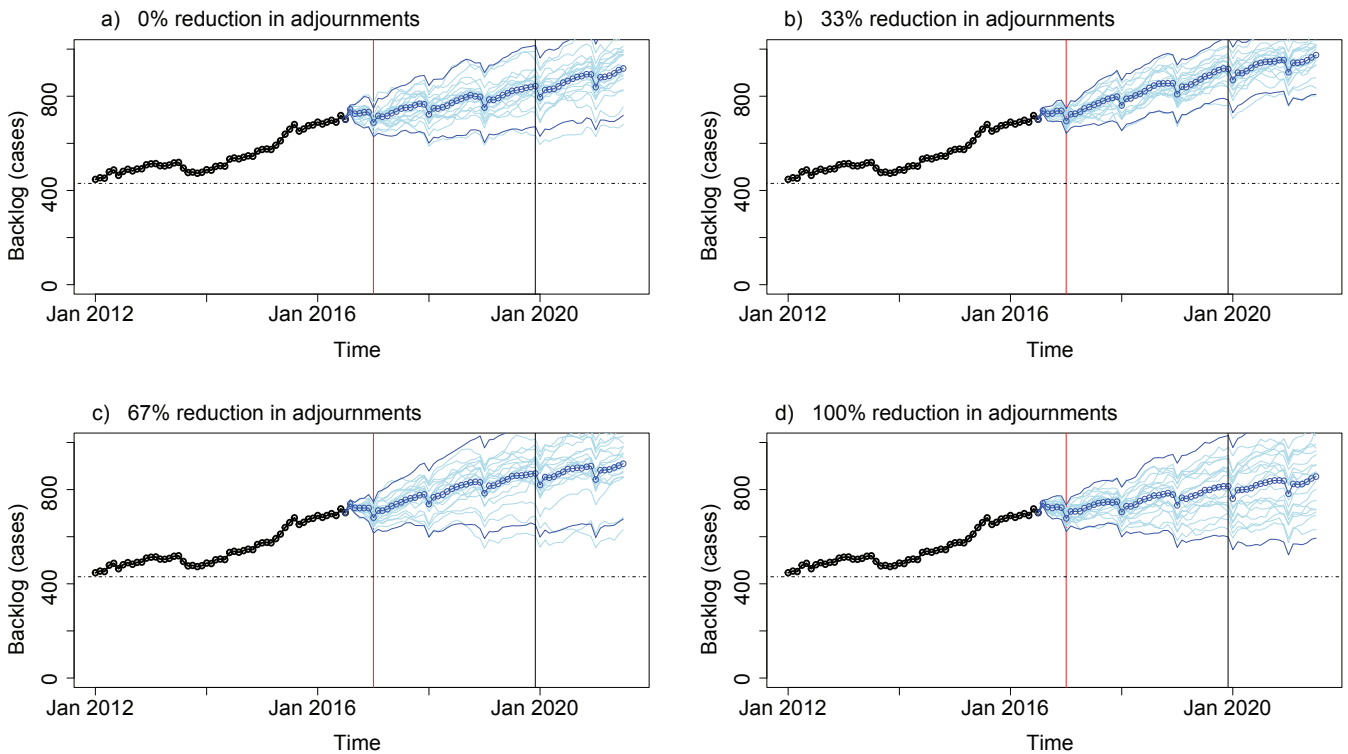


Figure 7. Size of case backlog vs. time for different reductions in the probability of adjournments for 20 simulation runs



(no reduction, a 25% reduction or a 50% reduction in late guilty pleas result in average backlogs of 796, 814 and 764 cases respectively). This shows that the effects from reducing the late guilty pleas are much weaker than the effects from adding additional judges, and that guilty plea reform would only slow the rate of increase in the size of the backlog, rather than cause a decline.

POSSIBLE INTERVENTION: REDUCTION IN ADJOURNMENT RATES

Given the suggestions made by Coopers & Lybrand WD Scott (1989) and Weatherburn & Baker (2000) associated with case adjournments, we examined the effect of reducing the likelihood of a case being adjourned on the morning of the trial. In the model, this translates to reducing p_a (and correspondingly increasing p_p , p_g , and p_o). The results are shown in Figure 7.

As can be seen, reducing adjournments by any amount has little effect on the case backlog. The average backlogs in each scenarios range from 762 (when adjournments are reduced by 100%) to 869 cases (when adjournments are reduced by 33%) by December 2019. While it may seem unusual that reducing adjournments by 33% actually causes the backlog to increase faster, this is probably just due to the randomness associated with the simulation runs.

POSSIBLE INTERVENTION: REDUCING VACATED COURT TIME

A reduction/removal of the mid-year vacation, as suggested by Coopers & Lybrand WD, was also simulated.

Figure 8 shows the effect of removing the court vacation (Figure 8b) vs the baseline level (Figure 8a). Removing the vacation causes the average backlog to drop to approximately 648 cases by December 31 2019.

COMBINING INTERVENTIONS

None of the interventions examined so far achieves the goal of a 430 pending cases by December 2019 except appointing five or more additional judges to the Sydney District Court. This would be quite expensive even without factoring in the additional costs associated with hiring more prosecution, defence and court administration staff and (potentially) building new courtrooms. It may be more cost-effective to make a number of smaller reforms to different areas criminal case processing. To test this idea, we ran our simulation making the following changes combined:

- The addition of two judges
- A 50% reduction in late guilty pleas
- A 33% reduction in adjournments

Given that reducing court vacations is quite a drastic measure, we consider implementing the other interventions both with and without removing the court vacation. The results of these measures are shown in Figure 9.

Comparison of the projected change in the backlog in Figures 9a and 9b shows that a combination of strategies can produce quite strong effects. If the court vacation is kept, but the other interventions are implemented, then the backlog will drop to 671 cases in December 2019. If the court vacation period is removed, then the backlog will drop to 564 cases by December 2019.

Figure 8. Size of case backlog vs. time both with and without vacated courts for 20 simulation runs

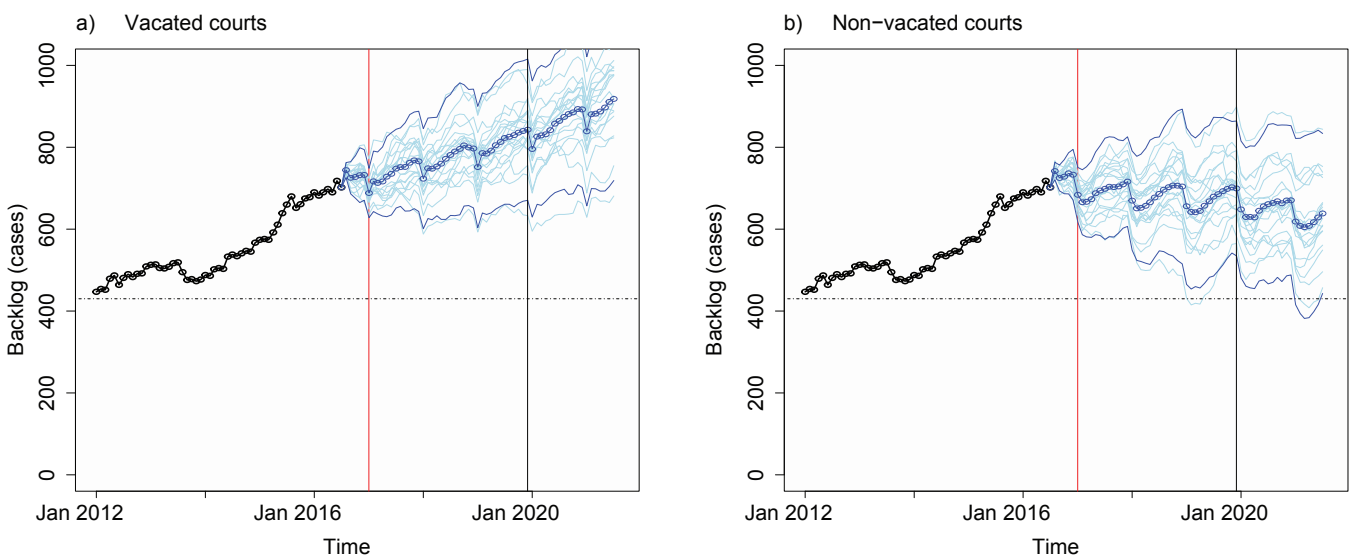
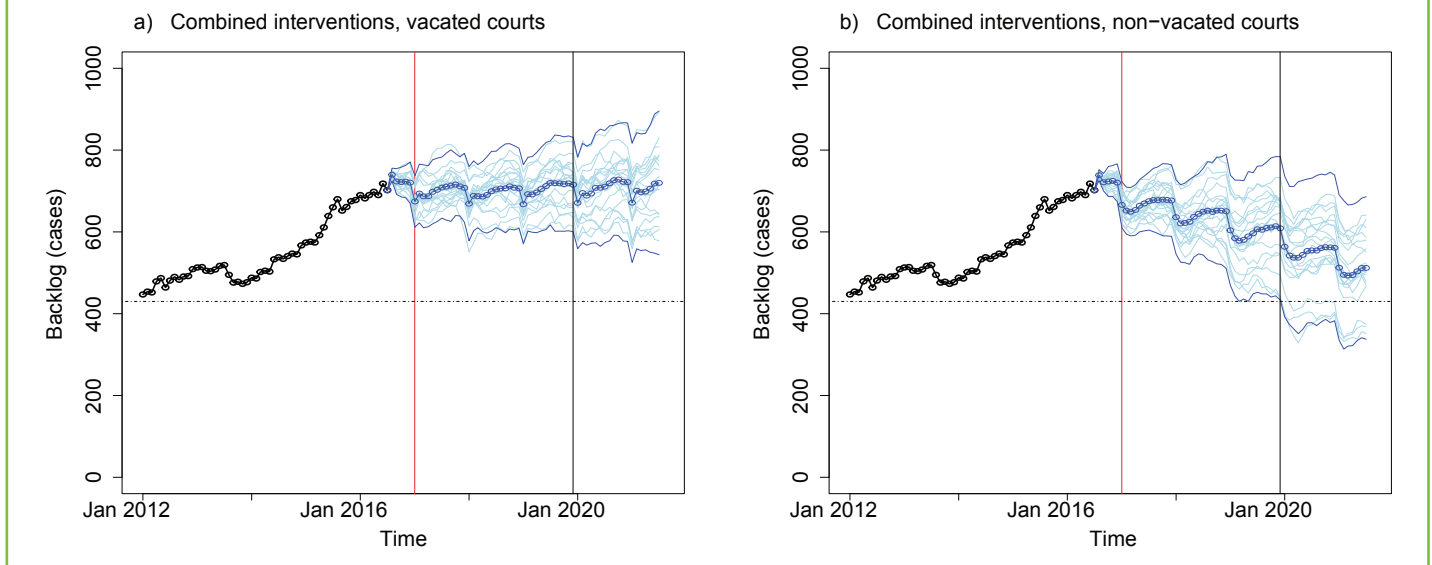


Figure 9. Size of case backlog vs. time for combined interventions both with and without vacated courts for 20 simulation runs



DISCUSSION

Given the importance of reducing the case backlog in NSW courts and the risks associated with implementing untested strategies for dealing with the problem, it is vital to have some means of assessing the likely effect of alternative strategies. The simulation model of the Sydney DCC reported here is designed to fill this void, allowing the user to simulate multiple possible reforms and interventions.

Looking at the results of the model, it is clear that the effect of additional judges has the greatest single impact. Indeed, the appointment of five or more judges is the only strategy which on its own would meet the target of 430 cases (in the backlog) by 2019. This is not to say it is the only strategy worth pursuing. As illustrated in Figure 9, interventions involving combinations of reforms/additional judges suggest that it is possible that fewer judges (e.g. two) could start reducing the backlog if combined with other reforms (guilty plea reform, reduction in adjournments and elimination of court vacation time). It is important to emphasize that these conclusions are based on assumptions about the impact of the policies that some readers might regard as unwarranted. In assessing the effect of early guilty plea reforms, for example, we assumed that they would have no effect on trial duration. If we were willing to assume that early guilty plea reforms would reduce trial duration, the estimated impact of early guilty plea reforms would be quite different to that reported here. In addition, our models have assumed that the interventions/reforms will be implemented instantaneously in January 2017. Clearly this is not realistic. However, these facts should not be taken as highlighting weaknesses of the

model; rather, they serve to underscore the point that good policy simulation in the context of court administration requires input from both technical experts and experts in court policy and court processes.

The importance of such cooperation is hard to overstate. To estimate the effect of a policy change, the user must decide which parameters to change (and to what degree they should be changed) to reflect the change in policy. While in some cases this is simple (e.g. a policy to introduce a number of new judges and courtrooms can easily be expressed as an increase in n), in other cases it is not. There are, for example, several ways in which one might choose to model changes to committal proceedings or new reforms to induce early guilty pleas. The estimated effect of policy change is also likely to be quite sensitive to the assumptions one makes about current trends in key parameters, such as the number of new trial registrations. Assessment of the sensitivity of simulation outcomes to changes in model assumptions should always form an integral part of the policy simulation process.

The model presented here can be improved in a number of ways. A problem which arose during the development of the model was difficulty in estimating the model parameter values. We only had access to court registration data from 2007 to 2016, which results in a very small sample size for the number of registrations each month for the different months of the year. Also, the data for trial-day outcomes and the proportion of committals reaching trial was difficult to estimate/verify between differing sources. In developing the model every effort was made to verify the chosen values of model parameters with experts

within the court system. The model predictions, however, remain quite sensitive to these parameters. Wherever possible, it would be beneficial to begin collecting systematic data on parameters whose true value is uncertain (e.g. the proportion of defendants who change plea on the morning of the trial)

Further extensions to the model are also possible. For example, it may be worth allowing model users to change the *distributions* (not only means) of the number of committals and trial duration (to help model changes to table 1/table 2 offences). Also, while the model currently outputs the size of the backlog, there are no outputs relating to 'time to justice' (the time between committal and finalisation). This would convert the model backlog into a measure more closely aligned with the courts' own service standards. Notwithstanding these limitations, the model is a powerful tool to be used by policy makers. It is hoped that beyond the results found here with regards to judges and early guilty pleas, the model could become a tool to be used by various parties in the Department of Justice when trying to determine future court reform.

ACKNOWLEDGEMENTS

The author would like to thank Jackie Fitzgerald, Suzanne Poynton, Nick Halloran (BOCSAR), Bill Hi, Brad Johnson, Rob Fornito (NSW District Court), Janine Lacey (ODPP), and Pip Hetherton, Stephen Bray, Alex Poulos and Malindi Sayle (Justice Strategy and Policy) for advice and guidance on the model structure and parameter estimation, Bill Hi, Brad Johnson, Rob Fornito (NSW District Court) and Janine Lacey (ODPP) for providing various data, external reviewers, Jackie Fitzgerald, and Don Weatherburn for draft feedback and Florence Sin for desktop publishing.

NOTES

1. Only the months of January to June were considered in when calculating the median court delay for 2016.
2. Here, the term 'vacation' refers to a period where judges do not preside over court proceedings. It is not necessarily the same as recreation leave (although judges often take their recreation leave during periods of court vacation).
3. A 'late' guilty plea is a guilty plea that occurs at any point after committal to trial in the District Court. A 'trial-day' guilty plea is any guilty plea that occurs on the morning of the listed trial.
4. A normally-distributed number can take any real value (i.e. the numbers will be decimal numbers). In our model, when a monthly number of registrations is randomly generated from a normal distribution, it is then rounded, as the number of monthly registrations must be a whole number.
5. Note that 'judges' in our model represent all required resources needed to hear a trial (judges, prosecutors, jury panels, courtrooms, etc.). Therefore, whenever this parameter is increased, one should be mindful also represents an increase in all other required resources as well.
6. The prediction intervals assume a normal distribution for the backlog at a given point in a simulation run. If the backlog is close to zero, this assumption may not hold, hence the prediction intervals will not be valid.

REFERENCES

- Chilvers, M. (2001) *Forecasting trial delay in the NSW District Criminal Court* (Bureau Brief No. 16). Retrieved August 24 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/BB/bb16.pdf>
- Crettenden, I., Packer, J. & Macalpine, S., (1993). *A computer simulation model of the District Criminal Court of New South Wales* (General Report Series No. 29). Retrieved 14 June 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/r29.pdf>
- Coopers & Lybrand WD Scott. (1989). *Report on a Review of the New South Wales Court System*. Sydney: NSW Department of Justice.
- Fornito, R. (2016). *Personal communication*, 19th of September, 2016.
- Greenberg, D. (1979). *Mathematical Criminology*. New Jersey: Rutgers University Press.
- Hetherton, P. (2016). *Personal communication*, 15th of September, 2016.
- Johnson, B. (2016). *Personal communication*, 2nd of September, 2016.
- Kuan, Y.Y., (2004). *Long-term trends in trial case processing in New South Wales* (Crime and Justice Bulletin No. 82). Retrieved 22 August 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/CJB/cjb82.pdf>
- Lacey, J. (2016). *Personal communication*, 5th of September, 2016.
- Lind, B., Chilvers, M. & Weatherburn, D. (2001). *Simulation the New South Wales criminal justice system: A stock and flow approach*. (General Report Series No. 50). Retrieved 22 August 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/r50.pdf>
- Livingston, M., Stewart, A., & Palk, G. (2006) *A micro-simulation model of the juvenile justice system in Queensland* (Trends

& issues in crime and criminal justice No. 307). Retrieved 14 September 2016 from Australian Institute of Criminology website: http://www.aic.gov.au/media_library/publications/tandi_pdf/tandi307.pdf

Nagel, S. S. (1977), *Modeling the Criminal Justice System*. Beverly Hills: Sage Publications Inc.

Needham, K. (2015, December 6). Baird government commits \$20m to help alleviate court delay 'crisis'. *The Sydney Morning Herald*. Retrieved 22 August 2016 from <http://www.smh.com.au/nsw/baird-government-commits-20m-to-speed-up-court-delay-crisis-20151205-glg8wp.html>

NSW Bureau of Crime Statistics and Research, (1989). Court delay and prison overcrowding (Crime and Justice Bulletin No. 6). Retrieved 14 September 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/CJB/cjb06.pdf>

NSW Bureau of Crime Statistics and Research. (2016). New South Wales custody statistics: Quarterly update June 2016. Retrieved 22 August 2016 from NSW Bureau of Crime Statistics and Research website: http://www.bocsar.nsw.gov.au/Documents/custody/NSW_Custody_Statistics_Jun2016.pdf

New South Wales Government, (2016a). *NSW Budget 2016-17 Media Release: Record investment in justice in NSW budget* [Press release] Retrieved 22 August 2016 from http://www.budget.nsw.gov.au/_data/assets/pdf_file/0007/128698/3a_Record_investment_in_Justice_in_NSW_Budget.pdf

New South Wales Government, (2016b). *The District Court of New South Wales Annual review 2015*, Retrieved 10 Sep. 2016 from: <http://www.districtcourt.justice.nsw.gov.au/Documents/2015%20District%20Court%20Annual%20Review.pdf>

Weatherburn, D., & Baker, J. (2000). *Managing trial court delay: An analysis of trial case processing in the NSW District Criminal Court* (General Report Series No. 47). Retrieved 13 September 2016 from NSW Bureau of Crime Statistics and Research website: <http://www.bocsar.nsw.gov.au/Documents/r47.pdf>

Weatherburn, D., & Fitzgerald, J. (2015). *Trial court delay and the NSW District Criminal Court* (Crime and Justice Bulletin No. 184). Retrieved 22 August 2016 from NSW Bureau of Crime Statistics and Research website: http://www.bocsar.nsw.gov.au/Documents/CJB/Report_2015_Court_Delay_cjb184.pdf

Worrall, J. (1982). A computer simulation of a court of summary jurisdiction. *The Australian & New Zealand Journal of Criminology*, 15(3), 154-162.

APPENDIX

APPENDIX 1: PARAMETER VALUES USED

The initial values used for the parameters in the model are displayed in tables A1 and A2.

Table A1. Parameters used in the model for mean and standard deviations of monthly registrations

Parameter	Month		
	Jan	Feb – Nov	Dec
A	25.20	56.48	48.50
σ	3.03	8.37	5.45

APPENDIX 2: INCREASES IN THE NUMBER OF COMMITTALS

In this section, we use an ARIMA model to determine the existence and magnitude of an increasing trend in monthly committals. The reason we use an ARIMA model is to account for the presence of autocorrelation and seasonality in the data for the monthly number of registrations. The monthly number of registered cases in the Sydney DCC between January 2012 and July 2016 is shown in figure A1.

Before fitting an ARIMA model, one needs to check whether the sequence of monthly numbers of registrations (referred to hereafter as the 'time series') is stationary. This was tested using a Dickey-Fuller test available in the STATA statistical software package. A deterministic trend term was also included. The test gave a *t* statistic of -6.495 which corresponds to a *p*-value less than .001. This indicates that the time series is stationary, so no differencing is required. The *t* value for the deterministic trend term was 0.4, corresponding to a *p*-value of .692, indicating that there is no deterministic trend present (although we are yet to account for seasonal differences).

The ARIMA model was fit using the 'forecast' package with R statistical software. This package is used as it contains the 'auto.arima' function, which selects the best ARIMA model (in terms of AIC) to fit to the data. Using this function, and including covariates for month indicators and a deterministic trend, it was determined that the best model to using was an ARIMA model with one seasonal autoregressive component (that is, a model with an autoregressive term at lag 12). The coefficients, standard errors, *t* statistics and *p*-values for this model are shown in table A3 (analysis was performed using SAS 9.4 statistical software).

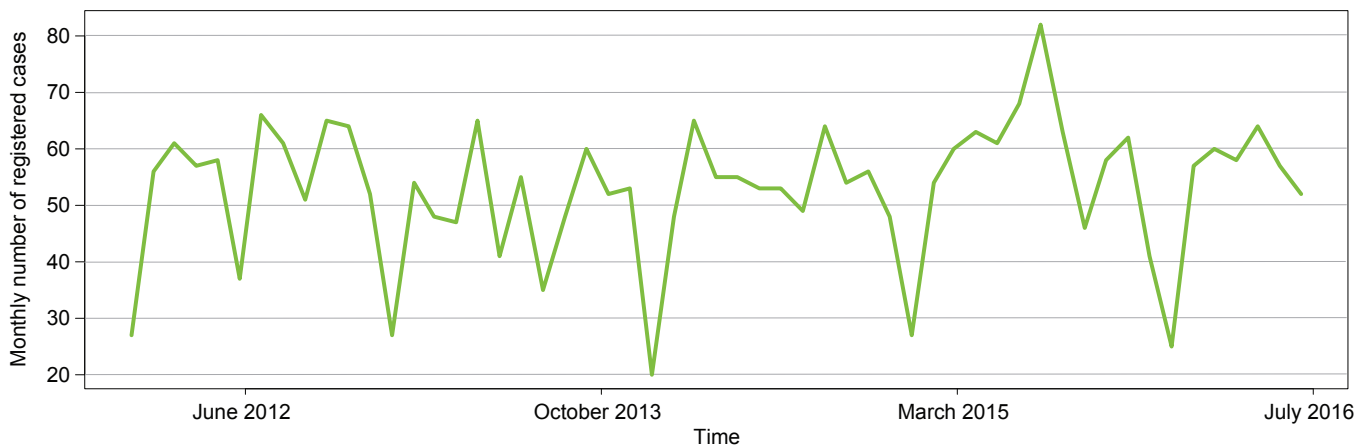
From Table A3, it can be seen that there no significant deterministic trend over time (*p* = 0.11).

For diagnostic checking, we use the Ljung-Box test to determine if any significant autocorrelations remain in the residuals of the ARIMA model (as this will violate our assumption for the ARIMA

Table A2. Parameters used in the model for trial day outcome probabilities

Parameter	Description	Value	Source/Validation
n	Number of judges used	18	User-set, but simulation runs with 18 judges and $p_r = .5$ match the observed pending caseload, and personal communication with the District Court
p_f	Probability that a case reaches the first day of a trial	.86	ODPP Data (Lacey, 2016)
p_p	Probability that a trial proceeds given it reaches the first day of the trial	.39	ODPP Data on finalisations combined with District court trial listing outcomes for adjournments
p_g	Probability that a defendant pleads guilty on the first day of the trial given he/she reaches day 1 of the trial	.34	ODPP Data on finalisations combined with District court trial listing outcomes for adjournments
p_j	Probability that a trial is adjourned on day one of the trial	.18	District court data on trial listing outcomes
p_o	Probability of a trial being disposed of by other means on the first day of the trial	.09	ODPP Data on finalisations combined with District court trial listing outcomes for adjournments
p_a	Probability of a trial being aborted/resulting in a hung jury	.05	District court data on trial listing outcomes
p_f	Probability that a judge can be reassigned a trial on the same days as an adjourned trial/trial that finalises as a LGP	.50	User-set, but simulation runs with 18 judges and $p_r = .5$ match the observed pending caseload, and personal communication with the District Court
M	Monthly increase in the number of registrations	0.0	District court registration data from January 2007 to June 2016.

Figure A1. Monthly number of registered cases for the Sydney DCC, January 2012 to July 2016



model). The results of this are shown in table A4. As can be seen, there are no significant autocorrelations remaining, hence the model assumption is satisfied and our ARIMA model in table A3 are adequate.

APPENDIX 3: TRIAL DAY PROBABILITY ESTIMATION AND RELATIONSHIP BETWEEN CHANGES TO LATE GUILTY PLEAS AND OTHER PARAMETERS

In this section, we consider how various trial day probabilities were calculated from the available data, and how certain parameters change when the rate of late guilty pleas is changed

by various amounts. The parameters we consider are p_g, p_f, A_m and M . We refer to the changed values of these parameters as p'_g, p'_f, A'_m and M' respectively.

Firstly, we show how we derive the probability of a case submitting a late guilty plea on the morning of a given trial, given the probability of a case being adjourned (which we obtained from District Court data), and the probability that a case is eventually finalised by a trial-day guilty plea (which we obtained from ODPP data). Let p_{fg} be the probability that a case which has reached the morning of the trial is eventually finalised via a

Table A3. Parameter estimates from the ARIMA model for monthly registrations

Variable	Estimate	Standard Error	t value	p-value
Intercept	17.11	5.63	3.04	.004
AR 12 term	-0.31	0.17	-1.83	.070
Deterministic trend	0.09	0.06	1.63	.110
Feb	28.31	3.88	7.30	<.001
Mar	33.32	3.88	8.59	<.001
Apr	30.43	3.88	7.84	<.001
May	35.01	3.88	9.01	<.001
Jun	25.85	3.89	6.64	<.001
Jul	36.21	3.90	9.29	<.001
Aug	25.59	4.15	6.17	<.001
Sep	27.39	4.14	6.61	<.001
Oct	33.62	4.14	8.12	<.001
Nov	32.47	4.15	7.83	<.001
Dec	23.20	4.15	5.59	<.001

Table A4. Results of the Ljung-Box test for autocorrelation

To Lag	Chi-Square	DF	p-value	Autocorrelations					
6	5.83	5	0.324	0.200	0.068	0.187	-0.074	-0.070	-0.086
12	11.12	11	0.433	-0.047	-0.227	0.045	0.120	-0.009	-0.089
18	15.91	17	0.530	-0.106	0.094	-0.091	0.056	-0.041	-0.163
24	29.98	23	0.150	-0.013	0.042	-0.102	-0.164	-0.121	-0.298

trial-day guilty plea. This is equal to the likelihood that the case is finalised immediately with a trial-day guilty plea, or that a case is adjourned and later finalised via a trial-day guilty plea (we will ignore hung juries/aborted trials here – given the likelihood of trials being aborted or resulting in hung juries, this effect should be negligible). This can be written as:

$$p_{fg} = p_g + p_a P(fg|a)$$

where $P(fg|a)$ is the probability that a case is eventually finalised via a trial-day guilty plea given that it was adjourned in that 'cycle'. However, we assume that the trial day outcome probabilities do not change across cycles – that is, a case which has already been adjourned is just as likely to plead guilty on the morning of the new trial as it was to plead guilty on the morning of its first trial listing. That is, $p_{fg} = P(fg|a)$. Subbing this into equation (1), we must have:

$$\begin{aligned} p_{fg} &= p_g + p_a p_{fg} \\ &= p_g + p_a(p_g + p_a p_{fg}) \\ &= p_g + p_a(p_g + p_a(p_g + p_a p_{fg})) \\ &= p_g + p_a(p_g + p_a(p_g + p_a(p_g + \dots))) \end{aligned}$$

Expanding out the parenthesis, we obtain

$$p_{fg} = p_g + p_a p_g + p_a^2 p_g + p_a^3 p_g + \dots \tag{2}$$

Equation (2) is a geometric series, with an initial term of p_g and a constant ratio term of p_a . We know that either $p_a = 1$ or $p_a < 1$. If $p_a = 1$, then the only way that equation (2) can be satisfied is by setting $p_g = p_{fg} = 0$, which makes intuitive sense – if cases are always adjourned, then they can never be finalised by a late guilty plea. As this is not the case, we must therefore have $p_a < 1$. In this case, the geometric series will converge and equation (2) can be shown to be equivalent to:

$$\begin{aligned} p_{fg} &= \frac{p_g}{1 - p_a} \\ \therefore p_g &= p_{fg}(1 - p_a) \end{aligned} \tag{3}$$

From data obtained from the ODPP, we estimate that $p_{fg} = .41$, and from the district court data, we estimate that $p_a = .16$. Therefore, we estimate $p_g = .41 * .84 = .34$. p_p and p_o can be defined analogously.

Reforms or interventions which reduce trial-day guilty pleas are treated as a reduced p_{fg} and correspond with an increase of guilty pleas occurring before committal in the Local Court. This means that A_m would be reduced, and (since proportionally

more cases that are committed are pleading guilty between committal and trial) p_f would also be reduced. The reduction in p_g is straightforward – since p_g is directly proportional to p_{fg} an $x\%$ decrease in p_{fg} results in an $x\%$ decrease to p_g . That is

$$p'_g = \frac{(100 - x)}{100} p_g \quad (4)$$

We now discuss the relationship between p_p and p'_p . Recall that p_{fg} is the proportion of committed cases which are finalised by a late guilty plea after reaching the morning of the trial, and let p_{fk} be the probability that a case is finalised by some other means after reaching the morning of the trial. Also, let p_{nf} be the probability that a trial does not proceed to its listing date, noting that $p_{fg} + p_{fk} + p_{nf} = 1$. Therefore, we must have

$$p_f = p_{fk} + p_{fg} = \frac{p_{fk} + p_{fg}}{p_{nf} + p_{fk} + p_{fg}} \quad (5)$$

Also, note that p_{nf} is made up of two parts – defendants who plead guilty at arraignment (a pre-trial hearing in the call to confirm the charges and set a trial date), and defendants who plead guilty between arraignment and trial. We will refer to these as p_{nfa} and p_{nft} respectively. Thus, Equation (5) becomes

$$= p_{fk} + p_{fg} = \frac{p_{fk} + p_{fg}}{p_{nfa} + p_{nft} + p_{fk} + p_{fg}}$$

If we reduce p_{fg} by $x\%$ (that is, $p'_{fg} = p_{fg} - \frac{x}{100} * p_{fg}$), then (5) becomes

$$p'_f = \frac{p_{fk} + p_{fg} - \frac{x}{100} * p_{fg}}{p_{nfa} + p_{nft} + p_{fk} + p_{fg} - \frac{x}{100} * p_{fg}}$$

As stated above, we assume that an $x\%$ reduction in late guilty pleas will result in all guilty pleas at arraignment now occurring at committal, and $x\%$ of guilty pleas occurring between arraignment and trial now occurring at committal. Therefore, we will have

$$p'_f = \frac{p_{fk} + p_{fg} - \frac{x}{100} * p_{fg}}{p_{nfa} + p_{nft} + p_{fk} + p_{fg} - \frac{x}{100} * p_{fg} - \frac{x}{100} * p_{nft} - \mathbb{I}_{\{x>0\}} * p_{nfa}}$$

$$= \frac{p_{fk} + p_{fg} - \frac{x}{100} * p_{fg}}{1 - \frac{x}{100} * p_{fg} - \frac{x}{100} * p_{nft} - \mathbb{I}_{\{x>0\}} * p_{nfa}}$$

where $\mathbb{I}_{\{x\}}$ is the indicator function, which takes the value of 1 whenever the condition in the braces is satisfied, and 0 otherwise.

Looking now at the relationship between A_m and A'_m . However, this is deceptively easy. Since we are reducing p_{fg} by $x\%$, we know that we must be reducing all finalised matters by $\frac{x}{100} p_{fg}$. Since each committed trial must be finalised, we know that the number of committals must reduce by the same amount.

Therefore, we must have:

$$A'_m = \left(1 - \frac{x}{100} p_{fg}\right) A_m$$

APPENDIX 4: VALIDATION OF SAME DISTRIBUTION FOR FEBRUARY TO NOVEMBER REGISTRATIONS

Table A5 shows the results for the regression of number of registered trials per month on month of registration

The estimate column of Table A5 shows the estimated mean for the referent month, as well as the relative difference for the mean of each other month. The p-values show the significance of the difference. As can be seen, only the p-values for December and January are less than .05, showing that only these two months are significantly different from April.

A Shapiro-Wilk test for normality was performed on the monthly number of registrations for the months of February to November. As the p-value for our test was .14, we can assume the data comes from a normal distribution. Due to the small number of observations for the months of December and January, it is not possible to validate the normality assumption for these months. While this may have an impact on individual simulation runs, this should not affect the calculation of averages over multiple simulation runs.

Table A5. Regression of number of registered trials per month on month of registration

Month	Estimate	Standard Error	t value	p-value
<i>Referent month (April)</i>	52.70			
August vs. April	-0.70	4.74	-0.15	.883
December vs. April	-10.81	4.74	-2.28	.025
February vs. April	-4.60	4.61	-1.00	.321
January vs. April	-26.10	4.61	-5.66	<.001
July vs. April	2.20	4.61	0.48	.634
June vs. April	-2.20	4.61	-0.48	.634
March vs. April	-1.30	4.61	-0.28	.779
May vs. April	2.60	4.61	0.56	.574
November vs. April	-1.26	4.74	-0.27	.792
October vs. April	1.97	4.74	0.42	.679
September vs. April	-1.14	4.74	-0.24	.810

Note: Months are listed in alphabetical order