Institute of Actuaries of Australia

An approach to improving claims management

Prepared by Bevan Damm and Daniel Marlay

Presented to the Institute of Actuaries of Australia
12th Accident Compensation Seminar
22-24 November 2009
Melbourne

This paper has been prepared for the Institute of Actuaries of Australia's (Institute) 12th Accident Compensation Seminar
The Institute Council wishes it to be understood that opinions put forward herein are not necessarily those of the Institute and the Council is not responsible for those opinions.

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: actuaries@actuaries.asn.au Website: www.actuaries.asn.au

Abstract

A key challenge in managing the claims process in a workers compensation scheme is improving return to work outcomes and reducing claims costs while also improving claimant health and social outcomes and improving customer service for all stakeholders in a scheme.

This paper takes a different approach to this issue to identify where there are opportunities to improve the claims management process.

In this paper we illustrate the approach by identifying where different claims managers are performing better than others, and then sharing that practice with the remaining claims managers to illustrate what the financial impact to the scheme could be.

For a workers compensation scheme, our approach shows savings of 15% of scheme claims costs and improvements in return to work rates of 3%.

The statistical techniques used group common *claims behaviours* together and measure the way the claims behaviours change over time.

Our analysis also identified significant differences between claims management by injury types; for lower back sprains there was a measured difference between agents of up to 34% of total claims costs.

Keywords: workers compensation; claims management; cluster analysis; state space

Introduction

Workers compensation schemes are complex due to the wide range of injuries, the volume and types of payments, the benefits structures and supporting legislation and the long life of claims. This makes the task of analysing a whole scheme while capturing the complexity and granularity of the individual payments difficult. There is also the difficulty of understanding the operations of the claims management and injury processes.

The challenge for all schemes and claim managers is to improve the outcomes both in terms of customer service and financially. One of the challenges in improving claims management is of course to understand what is currently performed and more importantly what "most effective" claims management really is. Once we know where we are and what needs to be done there is always the difficulty in measuring the performance.

In this paper we discuss an approach that can be used to identify the 'most effective' claims management practices from across a "universe" of practices. We have then modelled what the impact could be if those practices were adopted uniformly across a scheme. For this exercise we have used data from the NSW WorkCover scheme.

We don't attempt to say what 'most effective' claims management is, instead we assume that somewhere in the scheme it exists and we search through the data for it. To conduct this search, we group payment types together that represent certain types of claims behaviours and then project out the claim paths followed by these claims. To assess the 'effectiveness' of the claims management practices we have used two measures – total claims costs and a return to work measure rate.

We also recognise that the data we have used, while not perfect, is consistent enough to provide a demonstration of this approach. Naturally if this approach was to be used it would need to be further supported by discussions/claims audits with agents around their claims management practices in detail and the underlying data that is available.

This paper steps through the method we applied, the investigation and interpretation results; and provides the results of projections of the analysis. We draw conclusions regarding our analysis and projections then provide a list of further practical applications of this analysis.

Methodology

This section describes the technical detail of the analysis we have conducted. For a first read, the reader may skip ahead to the Investigation and Interpretation section and returned to later without detriment to understanding the conclusions of the paper.

Our approach uses common data summarisation and reduction techniques to simplify a complex data set and identify patterns of treatment. From these patterns we are able to investigate the trajectory of claim patterns for an injury across time. Across the portfolio these trajectories can be characterised in terms of transition rates between these claim patterns; we then compare these rates between agents. Finally, by adopting a state space approach, we are able to project outcomes for a portfolio of claims under different claims management assumption. This is summarised in Figure 1. For these process diagrams, the rectangles represent processes while the parallelograms represent data.

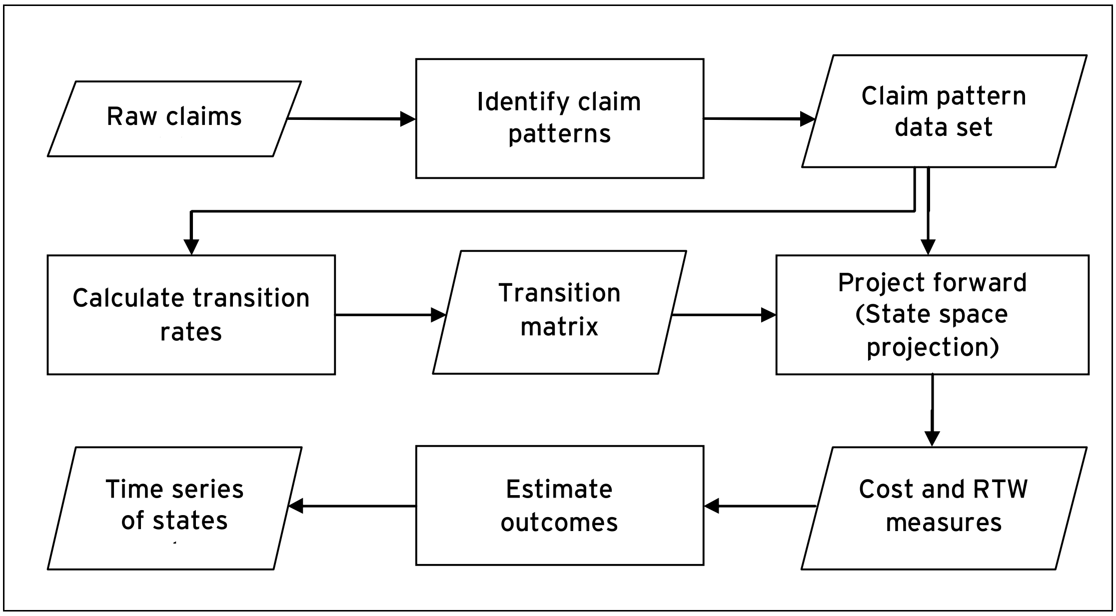

Figure 1: Overview of overall analysis and projection process

Identifying claim patterns

Figure 2 following provides an overview of the process of transforming the detailed claims data sets into characterised groups, called clusters. The cluster is used as a mechanism to provide a common description for underlying common types of claim payment behaviour. The objective with forming clusters is that the claim patterns in each cluster are distinctly different to those in other clusters. There is typically some trial and error involved in this process; as well as some individual claim patterns which do not conform well to the limited number of clusters formed. In general, while these outlier claims are interesting on their own, for the claims optimisation approach, the conclusions are identical due to the broad influence of the dominant claim payment patterns.

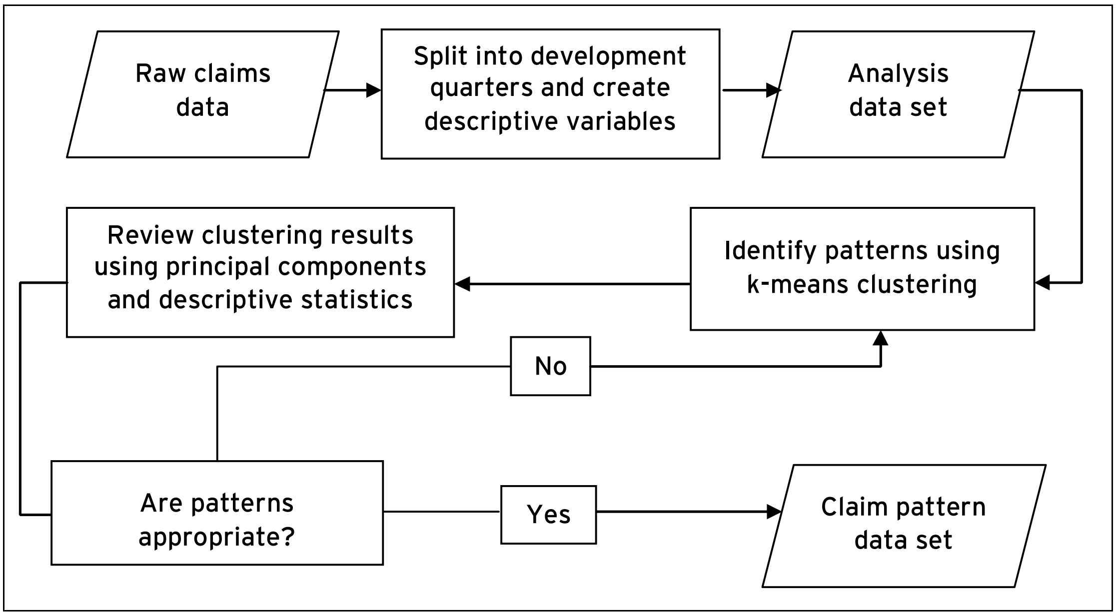

Figure 2: Overview of claim pattern identification process

A major challenge in understanding the behaviour of agents and claimants within a workers compensation scheme is in interpreting the volume of data. Each claim covered by the scheme will have many services for medical expenses, Weekly payments, legal costs and other items covered by the scheme. In order to understand the differences between claims, we need a way of summarising this data.

The first stage in our analysis is to break down the data associated with each claim into development quarters. This means that claims data for each claim is broken down into a series of quarter year periods, starting from the day of the claim. By breaking down the data in this way, we are able to investigate the claim patterns over time and understand how the treatment of the injury changes and how return to work outcomes are achieved.

To prepare the data for further analysis, the data associated with each development quarter for each accident is summarised into a single data point. This data point contains details for each type of claim (e.g. General Practitioner attendance, diagnostic imaging, Weekly payments, etc), the number of services received in the quarter, the total amount claimed and the average cost per service.

We have used a clustering technique to map the claim patterns for each accident in each development quarter to one of a number of archetypical claim patterns. This approach focuses on identifying development quarters that have similar claiming activity and mapping these quarters to the same archetype.

The clustering technique that we used in this analysis was the k-means clustering algorithm implemented by the SAS PROC FASTCLUS routine. This technique starts by selecting, at random, a fixed number (k) of archetypical development quarters from the full data set. The algorithm then repeats a process of assigning each development quarter within the data set to its closest archetype and updating the archetype to be the average of all of the development quarters assigned to it. This process stops when there is no change in the set of development quarters assigned to each archetype.

The k-means algorithm does not provide any guidance on the number of archetypes that exist within the data set and so this was determined through an iterative approach. Different numbers of archetypes were tried within the algorithm and a visual inspection performed to assess the goodness of fit for this number of archetypes.

The data set used in the clustering analysis has a large number of dimensions that can make visualisation difficult. We used a dimension reduction technique called principal components analysis to reduce the dimensionality of the data set. This allowed us to view the data set and assess the fit of the clustering analysis.

Principal component analysis provides an ordered set of dimensions – principal components – that are linear combinations of the original dimensions. These dimensions are ordered by the amount of variation the data set has across each dimension. This allows a small number of dimensions to be selected that capture a significant proportion of the variation within the data set. Viewing the data set as a scatter plot across pairs of these dimensions shows a significant amount of the structure within the data set.

Figure 3 shows a set of four clusters, highlighted by the four colours. The objective in forming the clusters is to form 'tight' clouds with distinct space between the clouds, so that in two dimensions, lines can be drawn between the clusters. Added complexity is created by requiring this separation to be conducted from multiple points of view, so the two dimensional lines become multi-dimensional geometric figures that encapsulate each cluster. For example, in Figure 3, Cluster 1 and 4 appear to be touching, but when viewed from a different direction, are predominantly separated.

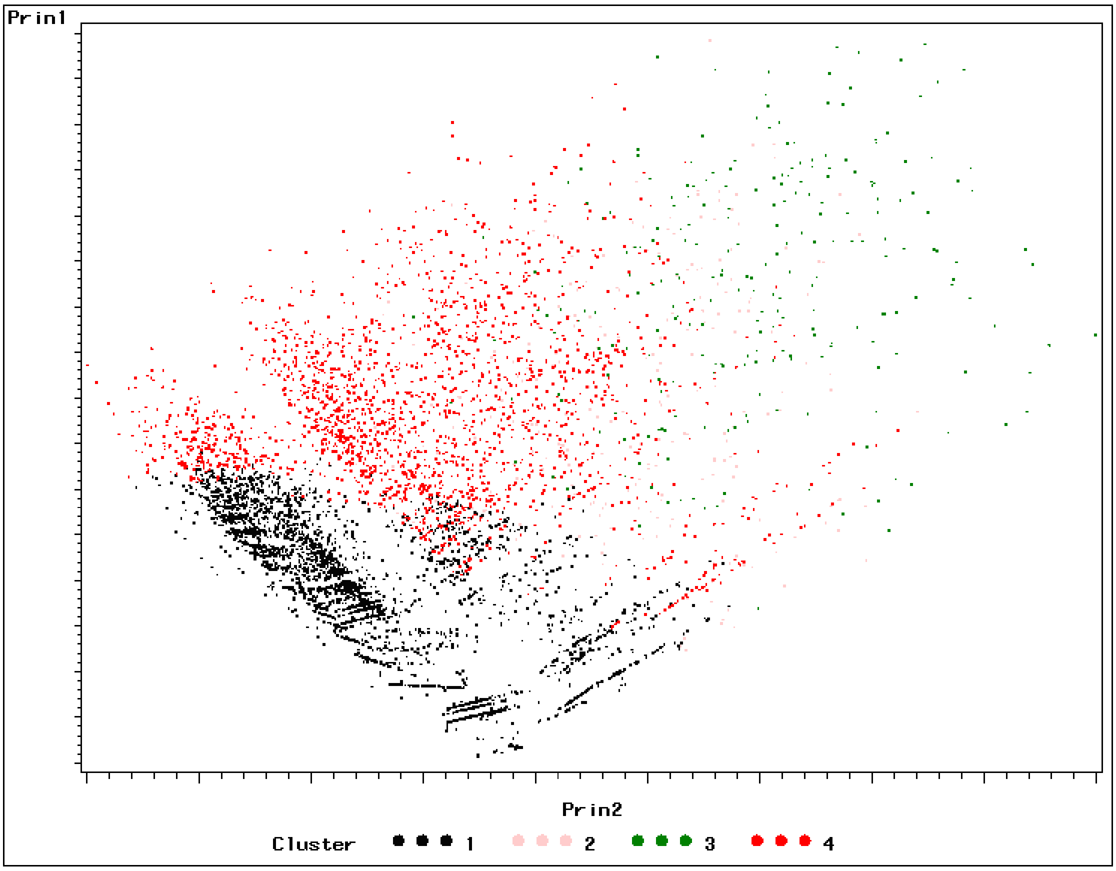

Figure 3: View of cluster clouds across the first and fourth principal component

Figure 4 following shows that Cluster 4 appears separated in both views, so may be a candidate for splitting into two distinct clusters. Other selected views of the final cluster clouds are provided in the Interpretation section.

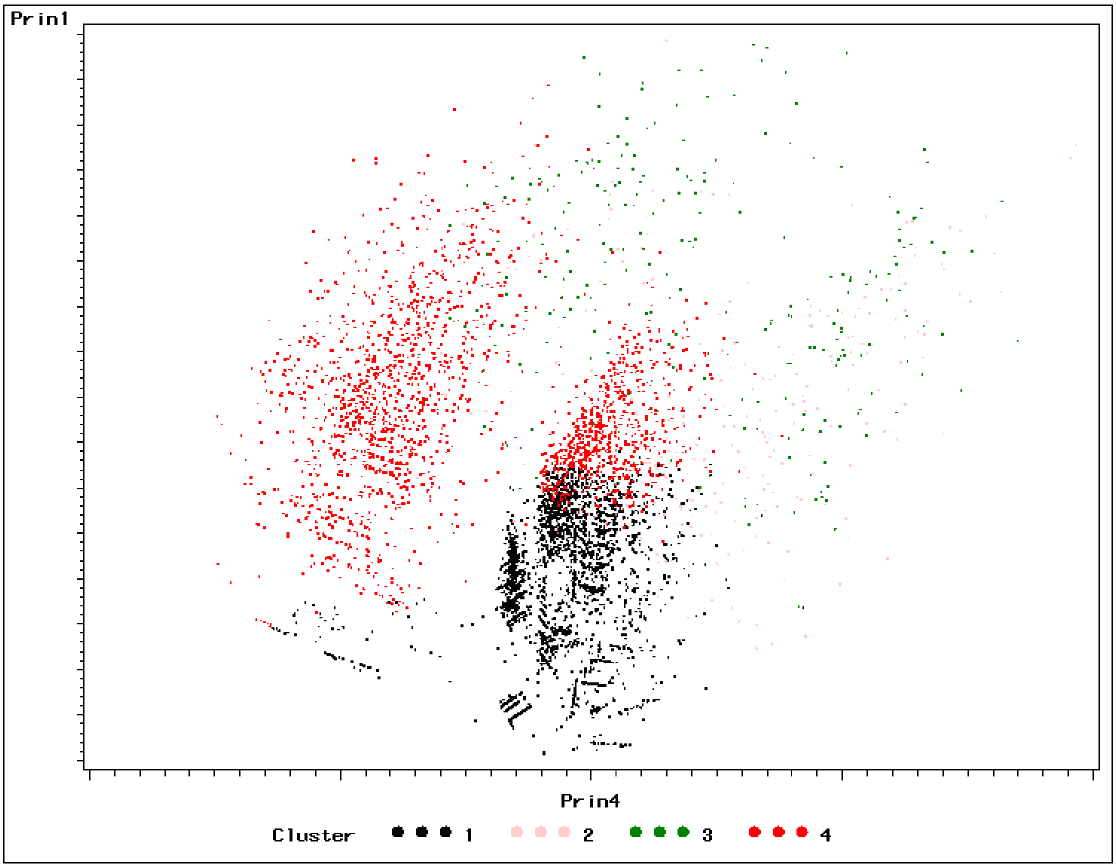

Figure 4: View of cluster clouds across the first and second principal component

Our analysis initially identified three key clusters; however two of these were associated with very atypical behaviour – death payments and lawsuits. We performed subsequent cluster analysis on the third cluster to further refine our characterisation of the development quarters resulting in five selected clusters that form the typical behaviour of claims payments in the scheme.

A final refinement was the creation of a final state cluster representing those development quarters in which an accident had no claim payments. While this is a short cut approach to claims finalisations, it also allows us to neatly capture re-opened claims.

We then examined the profile of each cluster across the different service types claimed. This allowed us to develop a qualitative characterisation of the cluster, to ensure that each cluster was generally distinct in its exhibited payment behaviour.

At this point in the process we have a claim pattern data set. This describes, for each accident for each development quarter for which we have complete data, the categorised claim pattern displayed in that development quarter.

Calculating transition rates

Figure 5 following provides an overview of the process of transforming the characterised claim payments into a transition matrix, which essentially provides typical 'claim paths' that claims will follow while active.

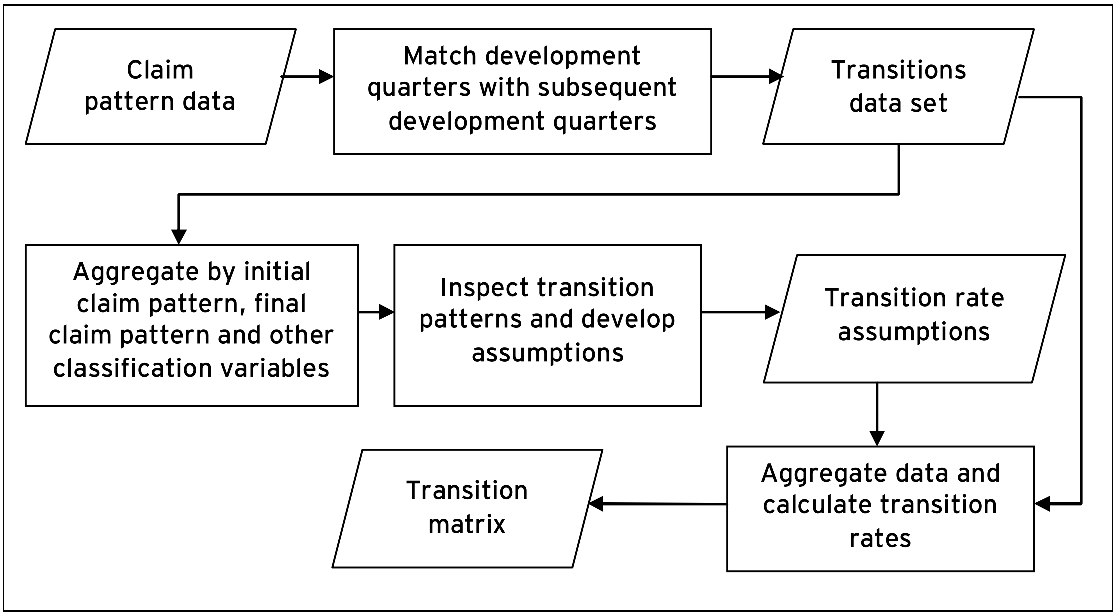

Figure 5: Overview of transition rate calculation

To understand the dynamics of a claim over time, we examined the transitions between claim patterns across development quarters. The result of this analysis is a set of transition matrices that can be used to perform the future projections.

The first stage in calculating the transition rates involves matching each development quarter of an accident with the subsequent development quarter. This produces a transitions data set containing, for each pair of successive development quarters, the initial claim pattern and the final claim pattern. In addition to data on the cluster, this data set contained details of the development quarter, the accident date, the agent and the injury type.

This data set was then aggregated by initial and final clusters to produce raw counts necessary for calculating the transition rates. The transition rates are calculated from the counts as:

$$r_{ij} = \frac{c_{ij}}{\sum_j c_{ij}}$$

where: r_{ij} is the transition rate from cluster i to cluster j

c_{ij} is the number of development quarters with cluster i that have a subsequent cluster j

These transition rates were calculated across different subsets of the data to examine the variation in transition rates over time (development and accident) as well as the variation between agent and injury type.

We assumed for the subsequent analysis, that the transition rates from one cluster to another were dependent only upon the agent, and not upon development time, injury type or calendar time. These assumptions are a simplification to allow subsequent analyses to progress, however a refinement to our methodology would either allow for these extra parameters or validate the simplified assumption.

Once our transition rate assumptions were fixed, we repeated the transition rate calculations to derive a set of transition matrices. These matrices were:

- An scheme transition matrix, with the average transition rates across all agents
- Separate transition matrices, with the transition rates for each agent, accident type and location

State space projection

Figure 6 provides an overview of the projection process which generates the time series of claim numbers that form the Return to Work and Claims Costs projections.

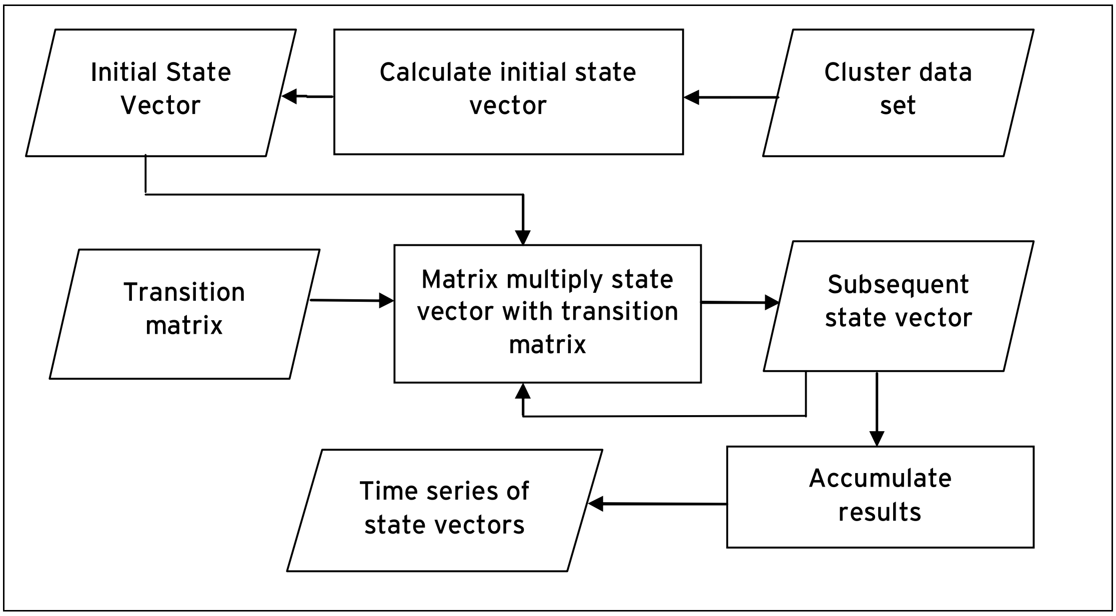

Figure 6: Overview of state space projection

The state space projection process takes an initial state vector, representing the current mix of claims, and uses a transition matrix to project this state vector forward in time. This allows us to simulate how a particular claims management approach, as represented by a transition matrix particular to an agent, impacts the outcomes for a portfolio of claims.

The initial state vector is calculated from the claims data set as the number of claims in each of the claim patterns, either for the entire scheme or for particular claim managers. This vector is calculated for both the entire portfolio and the sub-portfolio managed by each agent.

An estimate for each subsequent quarter's state vector was produced by multiplying the previous state vector by the transition matrix. This process, iterated starting from the initial state vector, was used to derive a time series of state vectors at quarterly intervals.

Our projection method made no allowance for the occurrence of new accidents, as our focus was on how a cohort of claims flows through the claims process differently for each agent.

The projection method allowed us to investigate a number of different scenarios, including:

- Projecting the current mix of claim patterns across the entire portfolio using the transition rates observed across the entire portfolio
- Projecting the current mix of claims for each agent using the transition rates observed for that agent
- Projecting the current mix of claim patterns across the entire portfolio using the transition rates observed with a particular claims handler. This is effectively simulating the effect of taking one claims management approach and applying it to the entire portfolio. This is the key to identifying which agent is operating the best

Estimating outcomes

Scheme management outcomes can be estimated from the state vector time series. In our approach we have calculated two measures of scheme effectiveness:

- Percentage of accidents that have been resolved ("return to work")

- Cost of claims for the portfolio of accidents

The percentage of accidents that have been resolved is estimated by assuming that accidents that have moved to the final state claim pattern have been effectively resolved. Thus a measure of the percentage of accidents resolved was calculated as simply the percentage of accidents that are in the final state claim pattern. It will be noted that not all accidents will resolve into the final state, and there is usually a tail of claims that remain in a different claim pattern receiving Weekly payments.

Cost is estimated from the average cost per quarter for each claim pattern. This is multiplied by the number of claims that are in each claim pattern and summed to provide an estimate of the total cost.

Assumptions

A number of considerations have required us to make use of simplifying assumptions in our analysis. Some key assumptions are:

- Claim management is static through time – in order to project the outcomes across the claims portfolios, we have assumed that claims management approaches (as measured using the transition rates) will be static into the future. This will not necessarily be the case; one of the desired outcomes from this work is to stimulate changes in claims management in order to improve outcomes for claimants and for schemes
- Transition rates are static with respect to development quarter and calendar quarter – our projections use a simple transition rate matrix. We assume that the likelihood of a claimant moving from one claim pattern to another remains constant, regardless of the time that has passed since the claimant had their accident or the calendar date on which the claim occurred

Investigation and interpretation

Claim patterns

Five clusters of claim patterns were identified in our cluster analysis. A summary of the number of claims, payments and services falling within each of these clusters is presented in Table 1 following. As noted in the method section, we had removed claim patterns associated mainly with death payments and lawsuits. These claim patterns were added back into the fourth cluster.

Table 1: Typical cluster key characteristics

Cluster	Number of claims #	Proportion of total claims %	Total payments $m	Proportion of total payments %	Number of services #	Proportion of total services %
Cluster 1	4,095	3%	44.0	15%	61,009	7%
Cluster 2	83,172	66%	144.9	49%	589,205	70%
Cluster 3	3,479	3%	42.0	14%	59,230	7%
Cluster 4	32,787	26%	47.0	16%	84,729	10%
Cluster 5	3,362	3%	16.7	6%	44,096	5%

Projections of these claim patterns across the first four principal components are shown in Figures 7a, 7b and 7c. A description of the cluster clouds was provided in the Methodology section.

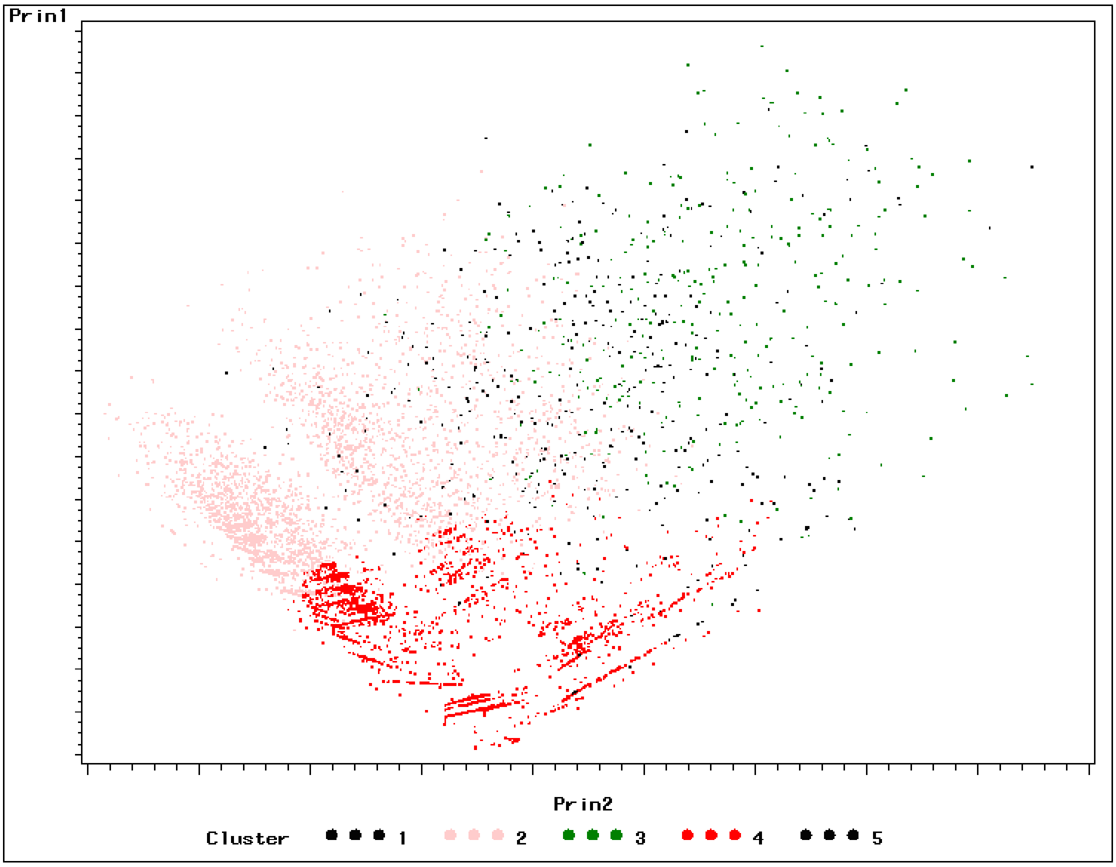

Figure 7a: Sample results of final cluster clouds

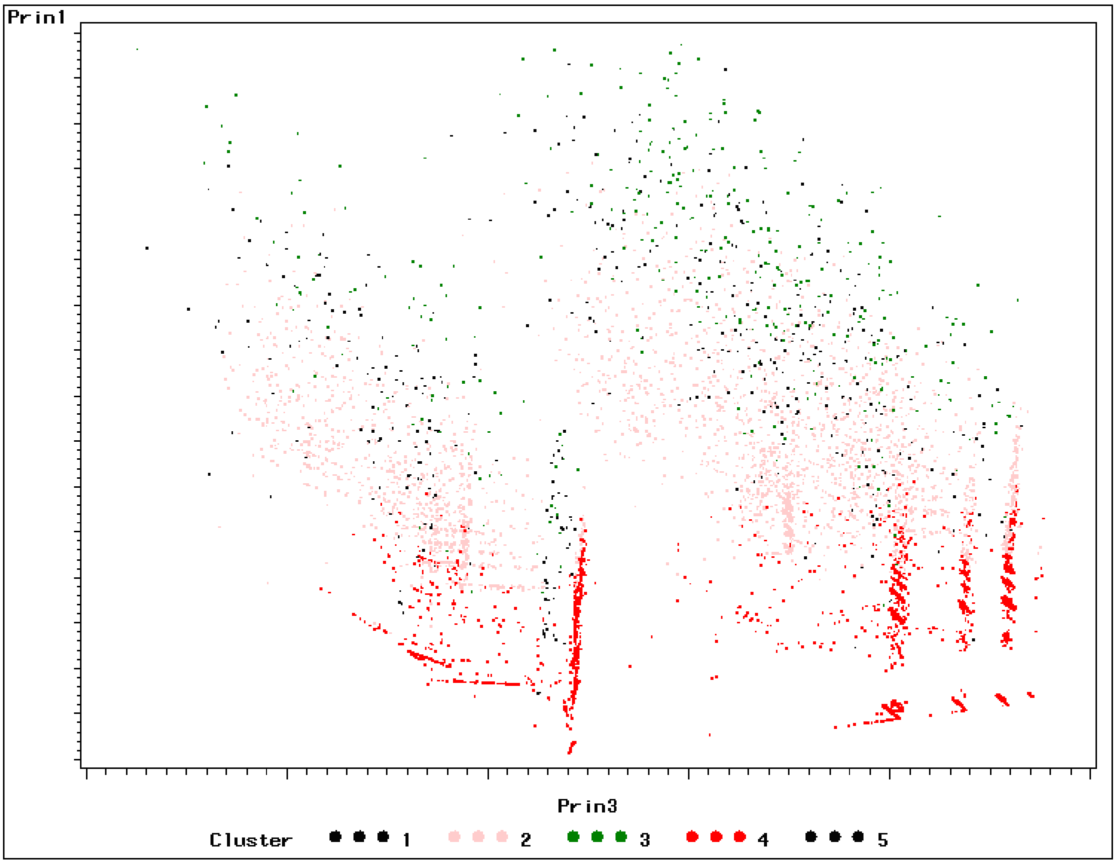

Figure 7b: Sample results of final cluster clouds

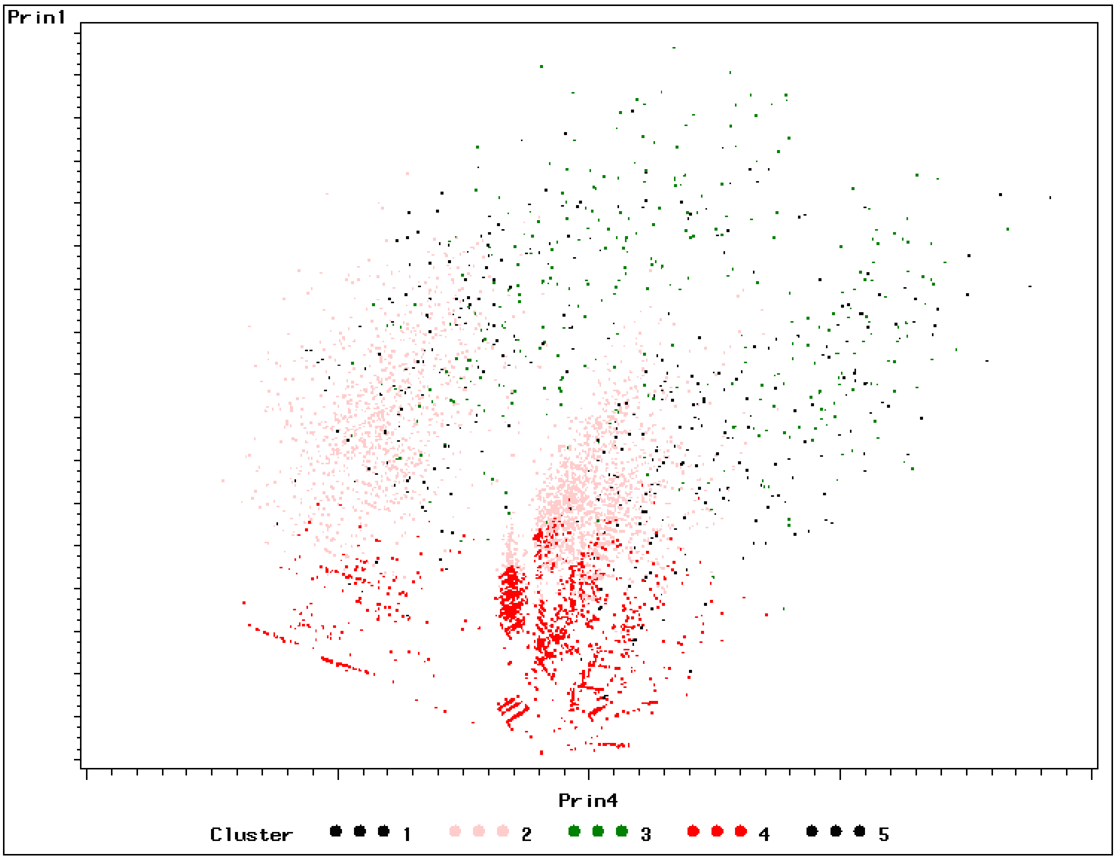

Figure 7c: Sample results of final cluster clouds

Table 2 following presents a summary of the key service and payment types represented within each of the clusters. This summary allows us to gain a better understanding of the type of treatment being received by a claim when it is in a particular cluster. Claims can receive service types not listed against their cluster, however they will typically receive smaller numbers of these services than those that are characteristic for their cluster.

Table 2: Typical cluster key characteristic payment types

Characteristic	Cluster 1 *Initial surgery*	Cluster 2 *RTW**	Cluster 3 *Further surgery*	Cluster 4 *Compensation*	Cluster 5 *Post-op recovery*
A	Surgery	Weekly	Weekly	Weekly	Weekly
B	Weekly	Occupational Rehabilitation	Private Hospital	Common Law	Orthopaedics
C	Public Hospital	General Practitioner	Orthopaedics	S66 and S67	Diagnostics
D	Private Hospital	Investigations	Surgery	Legal Services	Occupational Rehabilitation
E	Anaesthesiology		Anaesthesiology	Investigations	Physiotherapy

* *RTW* - payments that are representative of preparation for returning to work

The clusters in Table 2 each represent a typical claim 'treatment' pattern and represent current states of a claim. A label has been given to each cluster that represents the essence of the activity in that cluster. A summary of each of the clusters is described below:

- Cluster 1 *Initial surgery* includes ambulance costs, surgery and surgery related activities, and hospital treatment. Public Hospital is more prominent here than the other clusters due to the nature of the injury and the requirement of emergency treatment
- Cluster 2 *RTW* includes General Practitioner (GP) visits, Weekly payments and Occupational Rehabilitation. GP visits here are high because claimants are required to continue to show that they are in need of Weekly benefits in order for them to keep receiving them. This is more likely for patients receiving Care treatments and Occupational Rehabilitation until they can return to work
- Cluster 3 *Further surgery* includes Private Hospital, Surgery and surgery related activities. These represent non-immediate injuries that are treated after some investigative activities were undertaken or extra surgery after an initial emergency surgery
- Cluster 4 *Compensation* includes Weekly payments, Death payments and other Lump Sum payments. Claimants seeking Common Law, Commutations, Section 66 or Section 67 payments also received Weekly payments whilst not at work and required Legal Services and Investigation payments to justify payment of the Lump Sum
- Cluster 5 *Post-op recovery* is the smallest cluster dominated by Weekly payments, Diagnostics, Orthopaedics, Occupational Rehabilitation and Physiotherapy. This treatment pattern reflects claimants recovering from a surgical operation or claimants who are in between surgical operations

These clusters will become claim states for the projection process.

There are two other claim states required to represent the entering state (*Entry*) and the exit state (*Exit*):

- *Entry* reflects which state a newly injured worker first enters the scheme
- *Exit* reflects when claim payments cease. This generally means a claim is closed, however, it also allows for re-opening back into the active states

Table 3 following shows the typical cluster behaviour when a claim first enters the system, which is generally different to the behaviour once in the system. The main differences are in Cluster 2 *Light injury*, where GPs and Weekly payments dominate; and Cluster 4 *Initial injury*, where Weekly payments, GPs and Public Hospitals dominate.

Table 3: Entry cluster key characteristic payment types

Characteristic	Cluster 1 *Initial surgery*	Cluster 2 *Light injury*	Cluster 3 *Further surgery*	Cluster 4 *Initial injury*	Cluster 5 *Post-op recovery*
A	Professional Attendance	General Practitioner	Professional Attendance	Weekly	Orthopaedics
B	Surgery	Professional Medical Services	General Practitioner	General Practitioner	General Practitioner
C	Weekly	Weekly	Anaesthesiology	Public Hospital	Professional Medical Services
D	Anaesthesiology	Physiotherapy	Diagnostics	Physiotherapy	Diagnostics
E	Professional Medical Services	Diagnostics	Private Hospital	Professional Attendance	Professional Attendance

Table 4 summarises the claim states relative to the cluster, along with an interpretation of the claim patterns of the clusters that best explains the essence of the claim payments made while in that cluster. These descriptions are used in our subsequent analysis.

Table 4: Summary of claim state mapping and cluster descriptions

Claim State	Cluster	Description	
		First claim state	Later claim states
0		*Entry*	
1	Cluster 1	*Initial surgery*	*Initial surgery*
2	Cluster 2	*Light injury*	*RTW*
3	Cluster 3	*Further surgery*	*Further surgery*
4	Cluster 4	*Initial injury*	*Compensation*
5	Cluster 5	*Post-op recovery*	*Post-op recovery*
6			*Exit*

Transition rates

The combined scheme transition matrix is provided in Table 5. New cohorts of claims enter in the first column (*Entry*), where the majority enter into the *Light injury* cluster, with the majority of the remainder entering into the *Initial injury* cluster. The underlying behaviour of the clusters for the initial cohort of claims is different from the typical behaviour, as illustrated previously in Table 3. This indicates that most claim's first payments are for GPs, Public Hospitals, Weekly payments and Physiotherapy.

Table 5: Quarterly scheme transition matrix

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	3%	*Initial surgery*	12%	0%	3%	0%	1%	0%
Light injury	69%	*RTW*	19%	31%	34%	8%	28%	0%
Further surgery	2%	*Further surgery*	7%	3%	18%	1%	13%	0%
Initial injury	23%	*Compen -sation*	43%	32%	34%	63%	32%	3%
Post-op recovery	3%	*Post-op recovery*	3%	2%	9%	1%	16%	0%
Exit	0%	*Exit*	17%	32%	3%	27%	9%	97%

The overall scheme rates are shown in Table 5 above where the columns represent entry clusters and the rows represent exit clusters. The following interesting features can be observed:

- From *Entry*, a claimant is more likely to visit a GP or Public Hospital; and will typically receive Weekly payments and undergo medical testing
- Most claimants who were in *Initial surgery* were more likely to move to *Compensation*. A smaller proportion of claimants move to *RTW* and *Exit*
- From *RTW*, it is equally likely for a claimant to remain in *RTW* or move to *Compensation* or *Exit*; which means they will generally continue to receive Weekly payments and rehabilitation, receive a lump sum payment or leave
- From *Further surgery*, it is likely that the claimant will move to rehabilitation (*RTW*) or receive a lump sum payment (*Compensation*). They are also likely to require further surgery or physiotherapy (remain in *Further surgery*)
- From *Compensation*, it is likely that the claimant will continue to receive Weekly payments and eventually receive a Lump Sum payment and then exit from the scheme
- From *Post-op recovery*, it is quite likely for a claimant to move to Occupational Rehabilitation (*RTW*) or to a Lump Sum payment (*Compensation*). There is also a chance for the claimant to require further Surgery or more Physiotherapy; few claimants exit the scheme from this state
- From *Exit*, the majority of claimants remain out of the scheme however a few do return to receive Weekly payments or a Lump Sum payment

One of the key aspects of this work is to identify differences between claim managers, allowing investigation and transfer of the ‘best’ elements across all claim managers. Tables 6 to 8 provide a comparison of three claim managers’ transition matrices.

Points to note when comparing the agents’ rates:

- Agent C remains in *Initial surgery* more often than the other agents, although this may be related more to the injury than the agent
- Agent C exits at a lower rate than the scheme average and Agent A at a higher rate than the scheme average from the *RTW* cluster
- Agent A exits at a higher rate than the scheme average and Agent C at a lower rate than the scheme average from the *Compensation* cluster
- There are only small differences in transfer rates between agents for the *Further surgery* cluster
- Agent C tends to exit at a lower rate from *Post-op recovery*

Overall, Agent A tends to have higher than average exit rates while Agent C tends to have lower than average exit rates.

Table 6: Quarterly transition matrix for Agent A

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	3%	*Initial surgery*	11%	0%	2%	0%	1%	0%
Light injury	69%	*RTW*	18%	28%	35%	8%	29%	0%
Further surgery	2%	*Further surgery*	7%	3%	17%	1%	12%	0%
Initial injury	23%	*Compen -sation*	41%	28%	31%	56%	29%	2%
Post-op recovery	3%	*Post-op recovery*	3%	2%	11%	1%	17%	0%
Exit	0%	*Exit*	19%	38%	4%	35%	12%	97%

Table 7: Quarterly transition matrix for Agent B

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	2%	*Initial surgery*	13%	0%	3%	0%	1%	0%
Light injury	66%	*RTW*	19%	33%	33%	10%	28%	1%
Further surgery	2%	*Further surgery*	8%	3%	17%	1%	13%	0%
Initial injury	27%	*Compen -sation*	42%	33%	34%	61%	34%	4%
Post-op recovery	2%	*Post-op recovery*	2%	2%	8%	1%	16%	0%
Exit	0%	*Exit*	16%	28%	4%	26%	8%	96%

Table 8: Quarterly transition matrix for Agent C

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	4%	*Initial surgery*	17%	0%	3%	0%	1%	0%
Light injury	61%	*RTW*	17%	36%	30%	8%	24%	0%
Further surgery	3%	*Further surgery*	6%	4%	21%	1%	13%	0%
Initial injury	29%	*Compen -sation*	44%	37%	38%	73%	36%	4%
Post-op recovery	3%	*Post-op recovery*	3%	2%	6%	1%	21%	0%
Exit	0%	*Exit*	13%	21%	3%	17%	5%	96%

An issue of looking at the transition rates from agents in isolation, is that each agent is managing different mixes of claims. For example Agent A may be managing predominantly 'larger' employers while Agent C may be managing predominantly 'smaller' employers. To further isolate real differences between agents, one approach is to examine how each manager deals with the same type of injury.

Tables 9 and 10 following show two agent's transition rates for lower back sprain injuries. There are marked differences between the agents, with Agent A having considerably lower exit rates than Agent B, for the same injury type. Additionally, Agent A appears to have higher rates of transition from *Post-op recovery* back to *Further surgery*. This implies that it is worth investigating how each of these agents manages lower back sprain injuries to identify the key differences and transfer better practices between agents.

Table 9: Quarterly transition matrix for lower back sprain for Agent A

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	0%	*Initial surgery*	0%	0%	4%	0%	0%	0%
Light injury	78%	*RTW*	42%	24%	37%	20%	31%	2%
Further surgery	1%	*Further surgery*	8%	3%	19%	1%	21%	0%
Initial injury	21%	*Compen -sation*	50%	45%	31%	41%	35%	16%
Post-op recovery	0%	*Post-op recovery*	0%	2%	6%	1%	10%	0%
Exit	0%	*Exit*	0%	26%	2%	37%	3%	82%

Table 10: Quarterly transition matrix for lower back sprain for Agent B

From=> To	*Entry*	From=> To	*Initial surgery*	*RTW*	*Further surgery*	*Compen -sation*	*Post-op recovery*	*Exit*
Entry	0%	*Entry*	0%	0%	0%	0%	0%	0%
Initial surgery	0%	*Initial surgery*	20%	0%	1%	0%	2%	0%
Light injury	83%	*RTW*	50%	22%	52%	14%	44%	2%
Further surgery	0%	*Further surgery*	10%	2%	23%	1%	8%	0%
Initial injury	16%	*Compen -sation*	10%	36%	22%	34%	21%	11%
Post-op recovery	1%	*Post-op recovery*	0%	1%	2%	0%	11%	0%
Exit	0%	*Exit*	10%	39%	0%	52%	15%	87%

The final element required to finalise our search for optimal claims management is the average cost per quarter per cluster. Table 11 following provides quarterly average costs for each cluster for selected categories. Points to note are:

- Overall, the relativities between clusters is similar for most categorisations, with the exception of *Compensation* and *Post-op recovery* for Agent B when treating for a lower back sprain
- The average cost for *RTW* and *Compensation* is generally low, although claimants may be in these states for prolonged periods
- The average costs for interventions (*Further surgery* and *Post-op recovery*) is high, although these are generally only transitory states
- Lower back sprain generally attracts significantly higher ongoing treatment and compensation through the *RTW* and *Compensation* clusters; and tend to stay in these states for longer periods than average claims (shown as lower probabilities of leaving the relevant clusters in the transition matrices)

Table 11: Quarterly average cost of clusters for selected categories ($)

Cluster	Scheme	Agent A	Agent B	Agent A Lower Back sprain	Agent B Lower Back sprain
Initial surgery	10,753	10,911	10,825	11,879	13,467
RTW	1,742	1,807	1,919	5,818	9,490
Further surgery	12,076	13,797	11,927	21,734	19,936
Compensation	1,446	1,424	1,425	5,455	14,732
Post-op recovery	4,980	5,249	4,679	7,545	11,547

Projections

To further understand how the transition rates and average claim sizes of the clusters can affect the cost outcomes, stepped projections have been conducted. Starting with the scheme average of initial cluster mix, these projections compare several changing factors:

- Firstly, create a projection applying the scheme average claim size, using the different transition matrices for each agent
- Secondly, create a projection using each agent's average quarterly claim cost for each cluster as well as the individual agent transition matrices

The results are presented graphically over time in Figures 8 and 9 for lower back sprains. In Figure 8, two agents, with the same starting mix of claims (by cluster) are projected according to the relevant transition matrix. For this figure, since the cost per cluster is based on the scheme average, the difference between the lines reflects the different claim paths that each agent's claim management approach applies. The x-axis represents payment quarter, which is the same as development quarter in all of these projections, since the projections are of a single cohort of claims flowing through the scheme. The y-axis represents payments made per quarter, with the units for each projection in the figure caption. There is a marked difference through the claims path between the agents, with Agent A reaching a much lower equilibrium cost faster than Agent B.

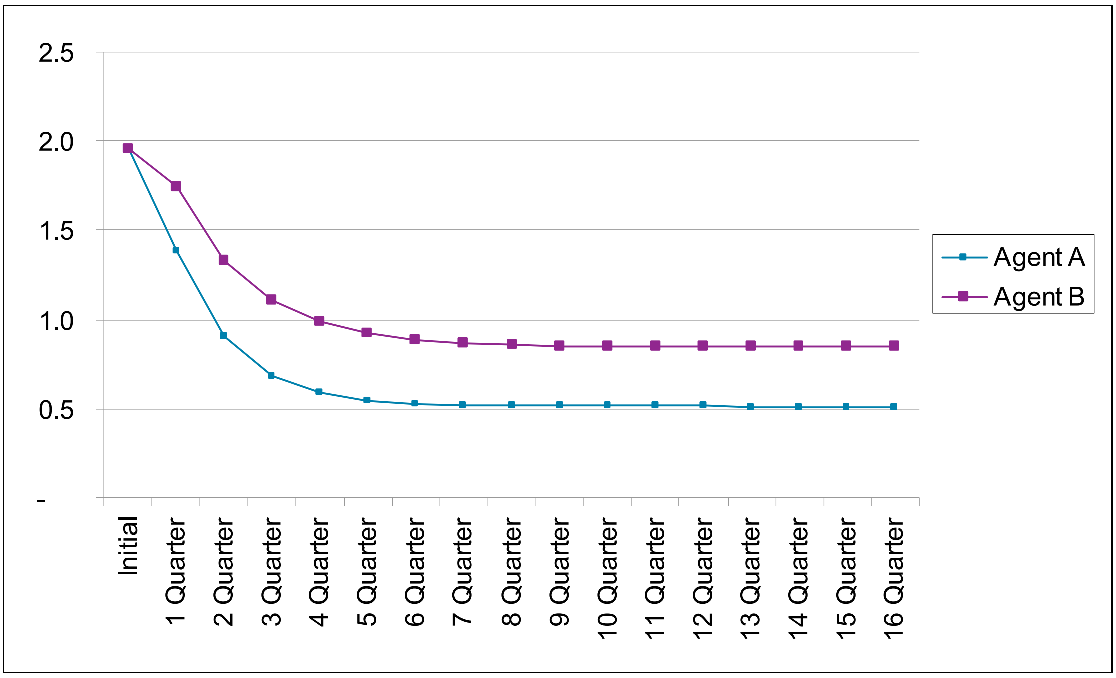

Figure 8: Lower back sprain – total claim costs per quarter; scheme initial entry mix and claim costs, individual agent transition rates ($m)

In Figure 9 following, the same claim paths are applied as for Figure 8, but here the average costs of the clusters specific to each agent are applied. This makes an even more significant difference to the overall costs of this type of injury between agents.

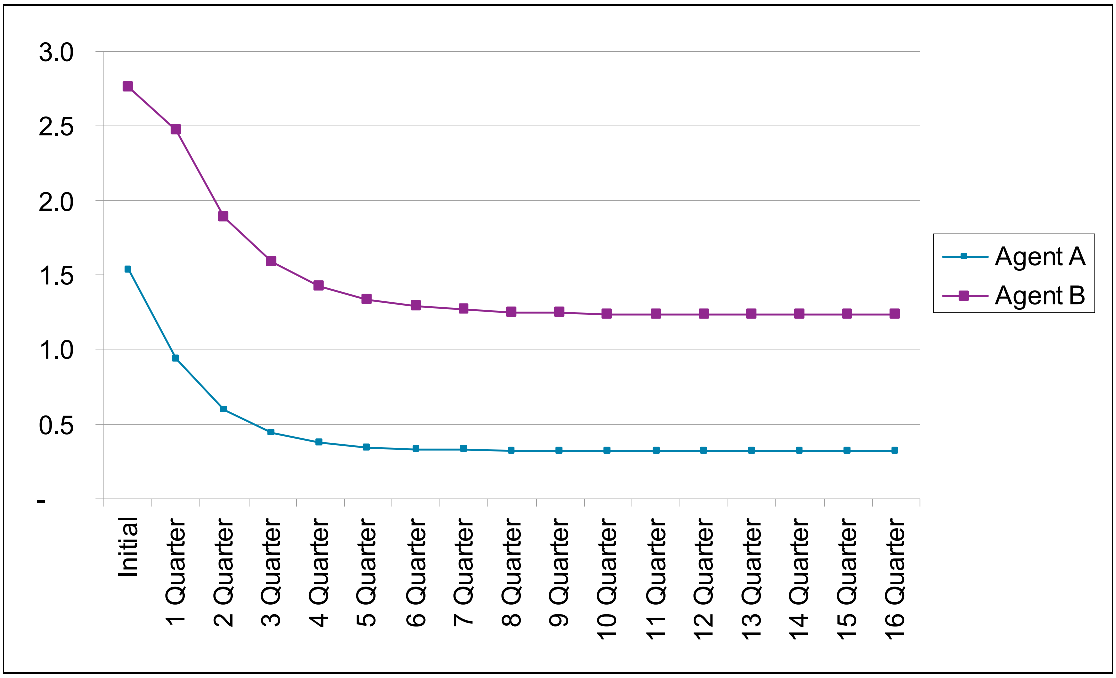

Figure 9: Lower back sprain - total claim costs per quarter; scheme initial entry mix, individual agent claim costs and transition rates ($m)

Two further injury types are shown, Finger open wounds in Figures 10 and 11, and Wrist fractures in Figures 12 and 13. Figure 10 illustrates that some injury types have very similar scheme outcomes across agents; whether that is through similar practices by the agents or by the providers is unclear.

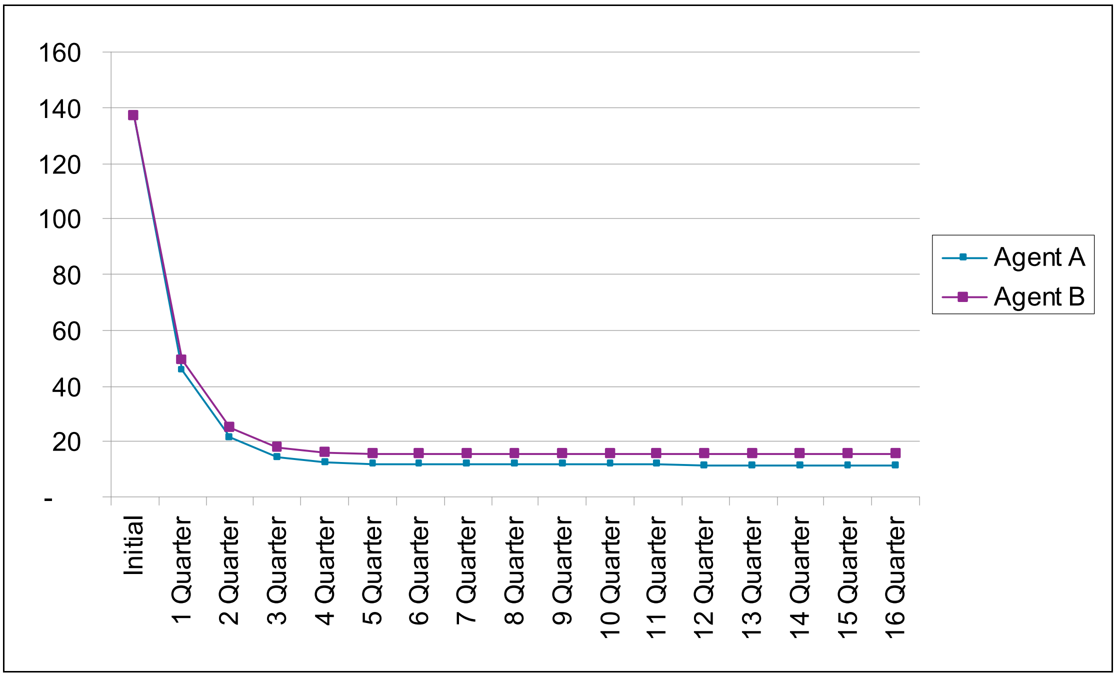

Figure 10: Finger open wound - total claim costs per quarter; scheme initial entry mix and claim costs, individual agent transition rates ($000)

Figure 11 following illustrates that, while there is minimal difference between agents in their claims paths, there is some difference created by the different costs structures of their respective clusters.

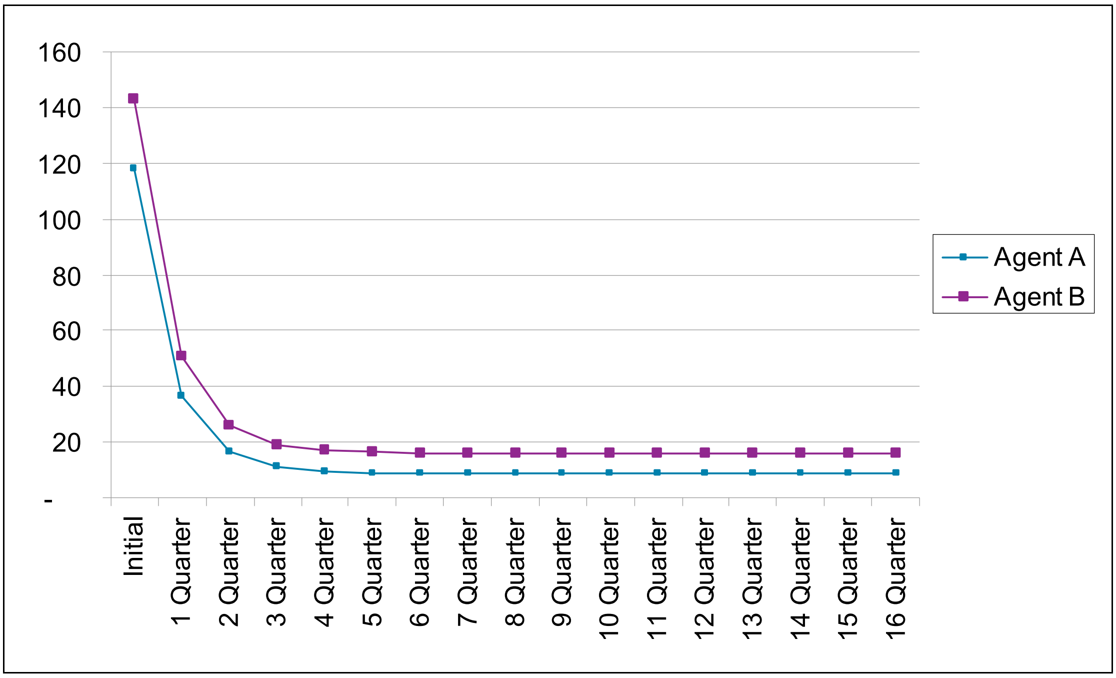

Figure 11: Finger open wound - total claim costs per quarter; scheme initial entry mix, individual agent claim costs and transition rates ($000)

Figures 12 and 13 illustrate that some injury types, while having different claim paths across agents, can also have average cluster costs that negate some differences between agents. Figure 12 shows Agent A has a better claims path, but Figure 13 shows that Agent A also appears to have a higher cost of following that path (the gap between long term costs has reduced in Figure 13 compared to Figure 12).

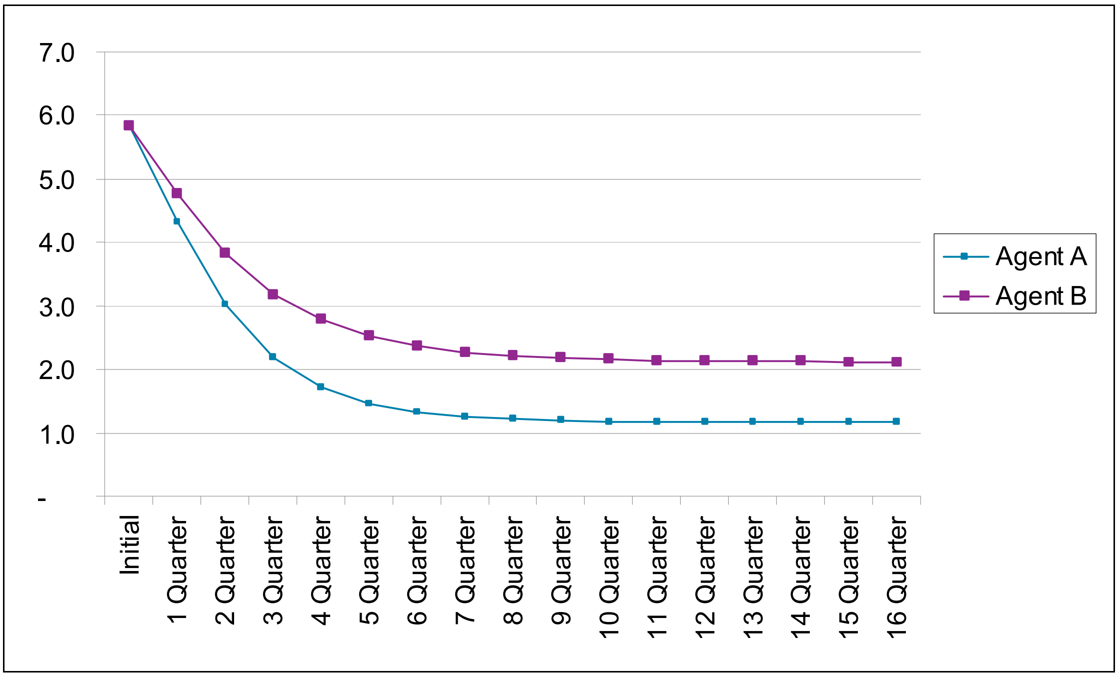

Figure 12: Wrist fracture - total claim costs per quarter; scheme initial entry mix and claim costs, individual agent transition rates ($000)

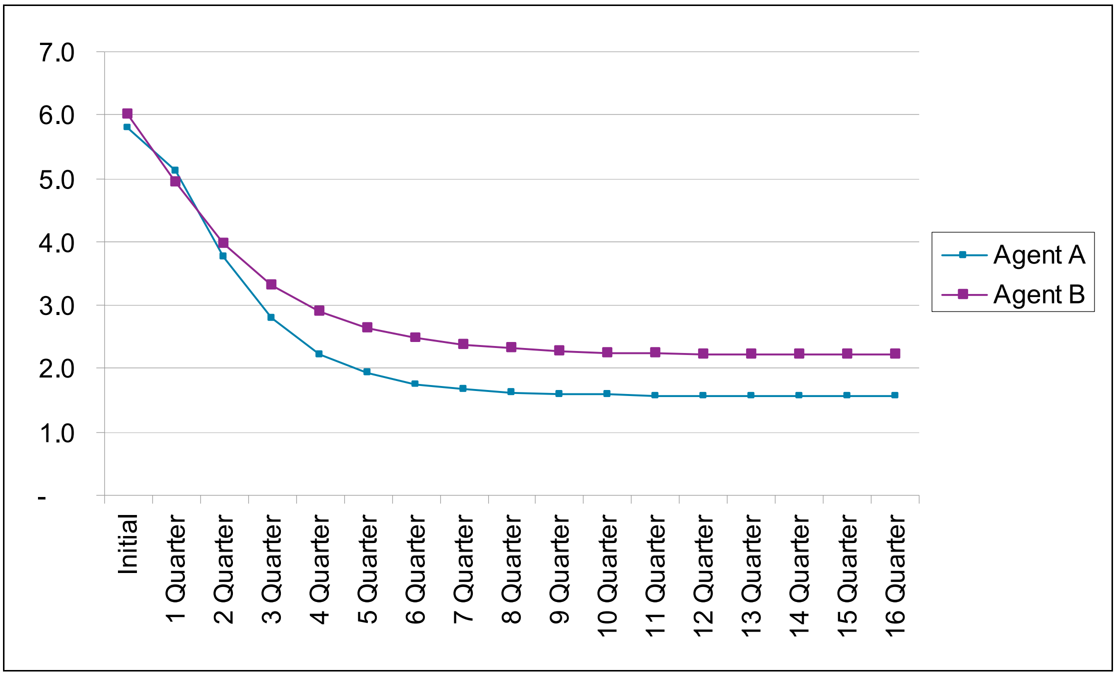

Figure 13: Wrist fracture - total claim costs per quarter; scheme initial entry mix, individual agent claim costs and transition rates ($000)

There are two approaches to optimising the claims management process that we have considered. Firstly, a Return to Work ("RTW") measure, where we are interested in the claimant being able to return back to work in the shortest time. Secondly, an overall cost of claims approach, where we are interested in minimising the overall costs to the scheme of a particular claim, which may allow a slower return to work if the cost structure for the longer active claim time is lower.

Generally, the two objectives would be considered to be complementary, although if the return to work occurs too quickly, then there may be excessive re-opening of claims and the service provided to the claimant may be lower than desired. One could assume that there is some mixing of the two objectives that would allow the shortest times and lowest overall costs.

Figure 14 following shows, using the overall scheme transition rates and cluster costs, the average RTW progress and claims costs over successive development quarters. The RTW rates show that almost all claims have achieved a return to work outcome within 2 years, with the majority of the higher expenditure rate from the high cost clusters also occurring during that time. Once equilibrium is reached the RTW rate remains static at 90%.

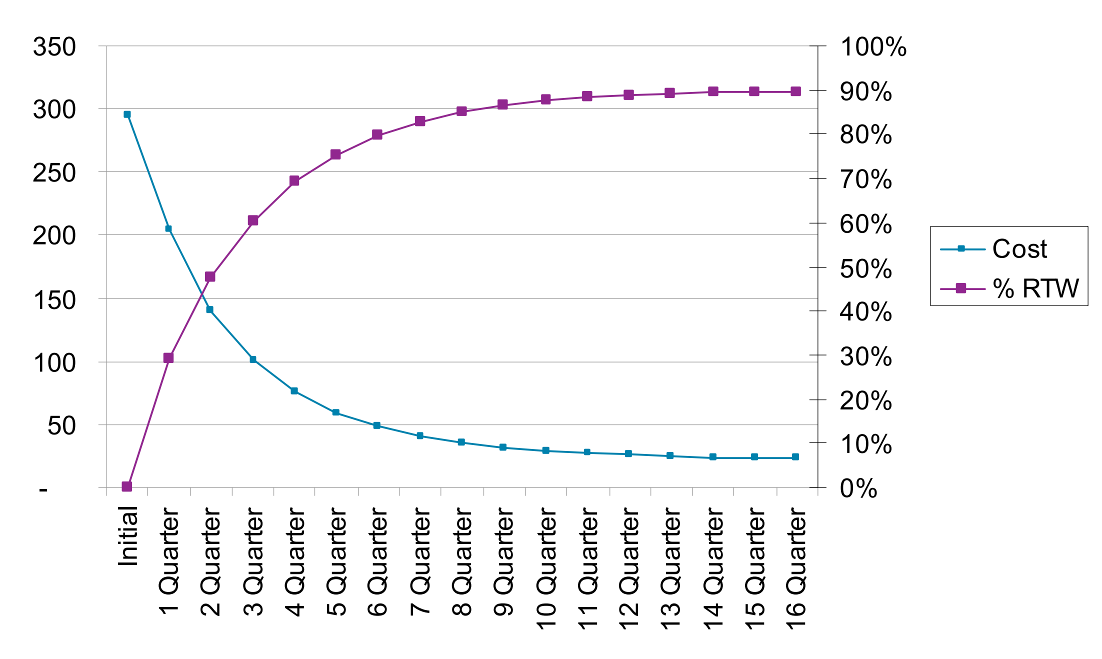

Figure 14: Scheme RTW and total claim costs per quarter ($m)

From the cost projection we can also create an equivalent present value of the projected claim costs, which can then be used to compare different approaches. The base projection present value is $1.2b. This represents one single accident quarter of new claims flowing through the claims process over time. To represent the entire scheme, an initial cohort needs to be applied every quarter as each new accident quarter of claims enters the system.

Figure 15 shows a breakdown of the average path of claims for the scheme. From this it is clear that after two years, equilibrium is approached, with the residual active claims remaining in the *Compensation* cluster. This cluster is dominated by Weekly payments, with the residual payments based on other investigation and compensation.

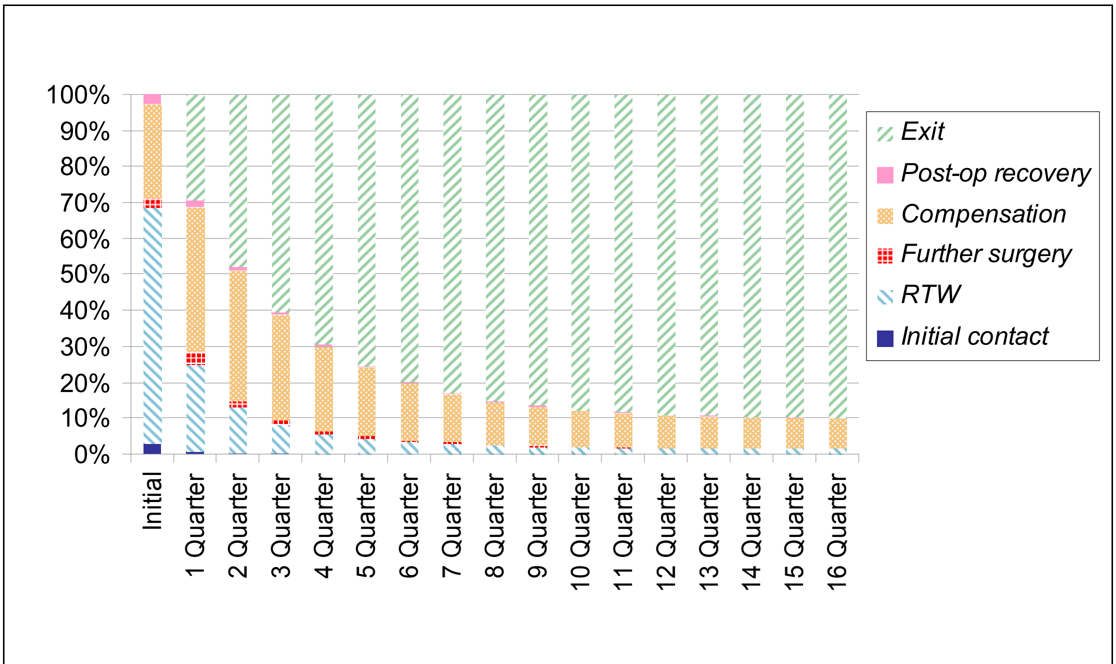

Figure 15: Breakdown of scheme cluster distribution over time

If we examine the RTW measure, this is effectively a measure of how quickly the claim reaches the *Exit* cluster, so the claim path that is followed through the transition matrix becomes important. Some paths, as has been shown previously, have higher costs than others. However, trade offs may occur when the claims path is through a higher cost cluster for a shorter time with faster subsequent recovery (i.e. more intensive targeted treatment) producing a faster return to work for a similar overall cost. Currently, we see some of this optimum claims management together with long term relatively low cost claims paths. The draw back is that these longer term claim paths tend to also use more claims management resources.

If we examine how two different agents differ over time, we find that the better RTW rates are generally aligned with lower overall claim costs. Figures 16 and 17 following show two agent's projections. Agent A is achieving RTW rates of 93% while Agent B is only achieving rates of 80% and achieving those rates much more slowly than Agent A. This assumes that they each start with the same mix of claims, so the difference between the RTW rates is based on each agent's claims paths that underlie their respective claims management processes. Similarly, when cluster costs are overlaid, the cost differentials are significant, with Agent B retaining a residual quarterly payment more than twice that of Agent A. The present value comparison, equivalent to the overall cost of claims management is $1.0b and $1.7b for Agent A and B respectively.

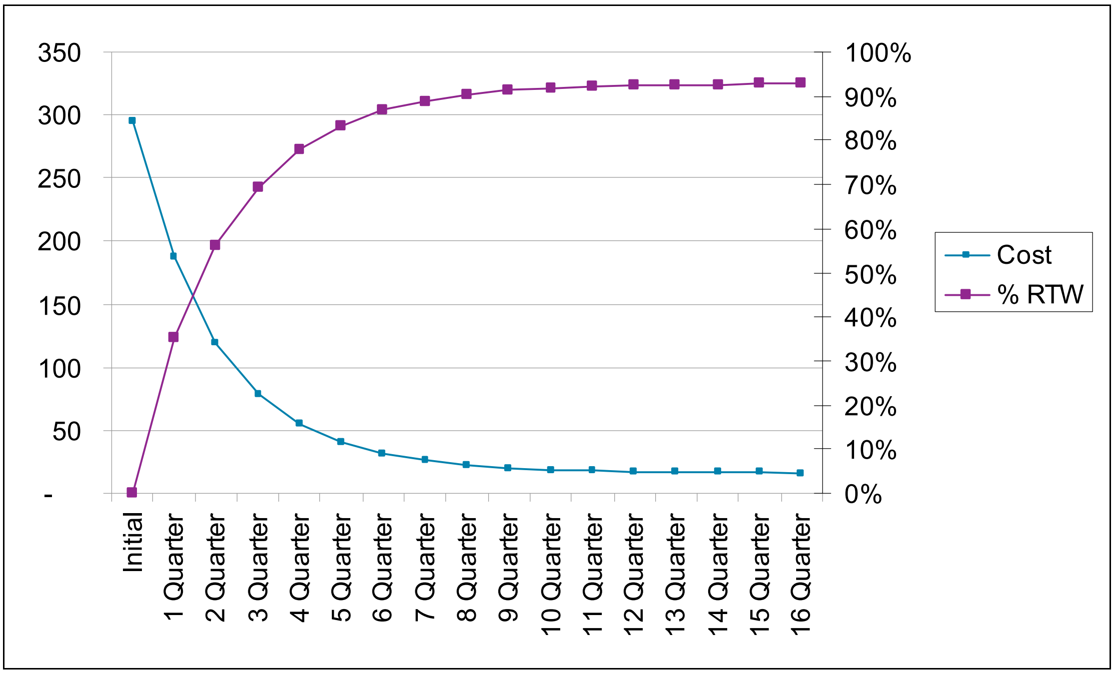

Figure 16: Agent A RTW and total claim costs per quarter ($m)

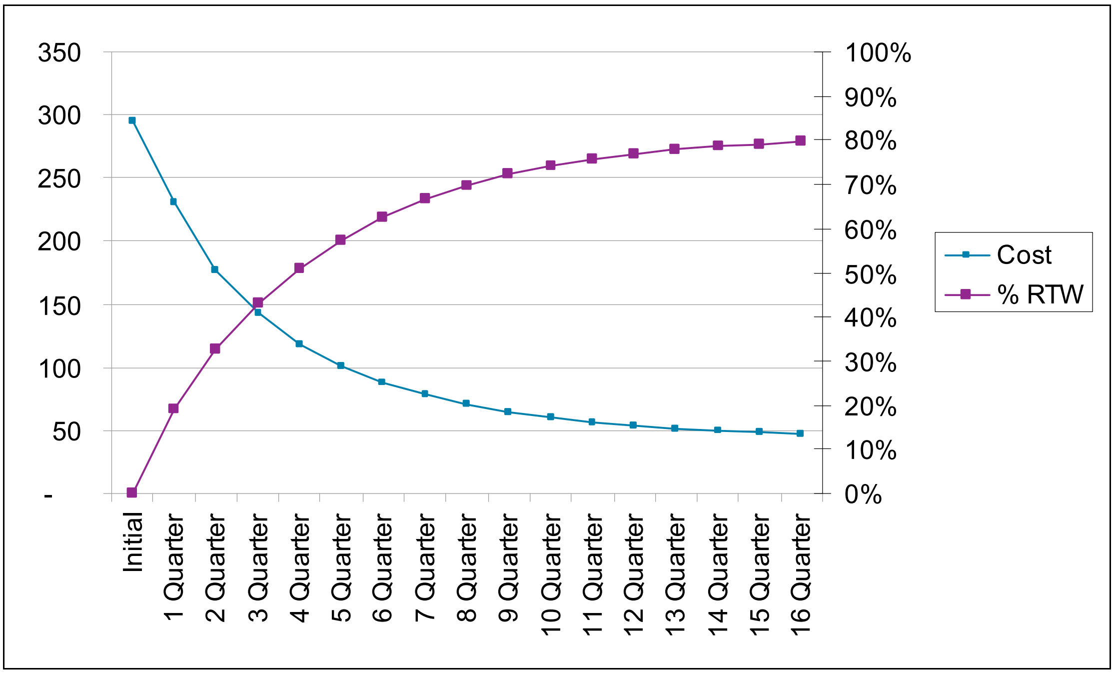

Figure 17: Agent B RTW and total claim costs per quarter ($m)

Figures 18 and 19 following show the breakdown of the cluster distributions for each of the two agent's projections. Agent A is minimising the *RTW* cluster and retains only *Compensation* as the residual cluster in the long term. Agent B is much slower in progressing claims through the claims path and after 16 quarters still has claimants in the *RTW* cluster and a significant proportion in the *Compensation* cluster.

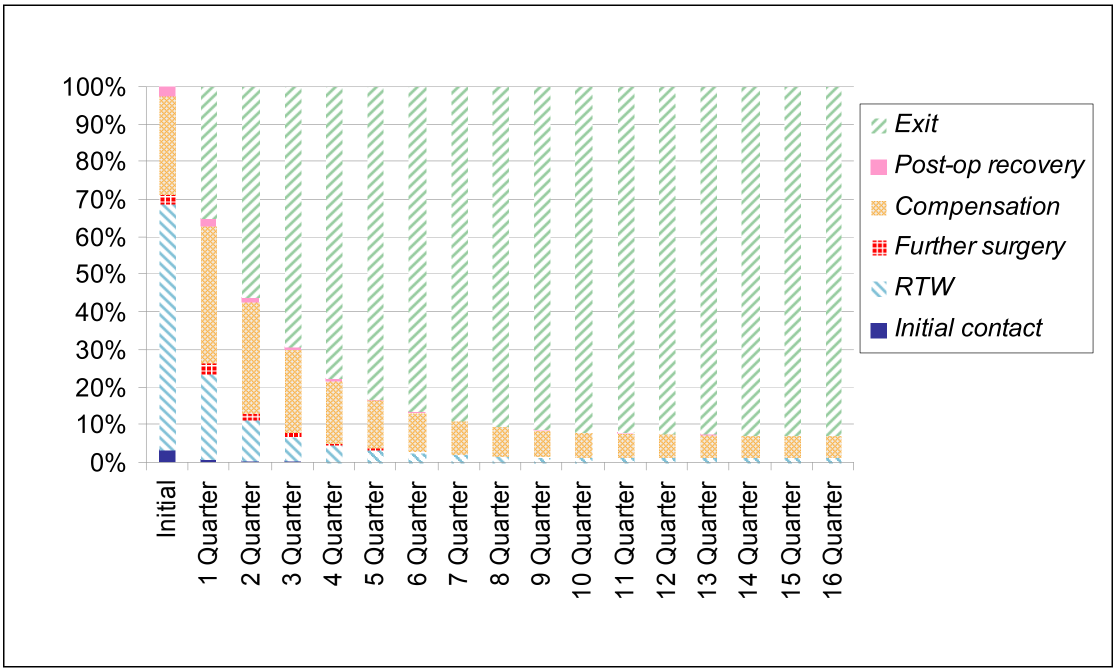

Figure 18: Breakdown of Agent A cluster distribution over time

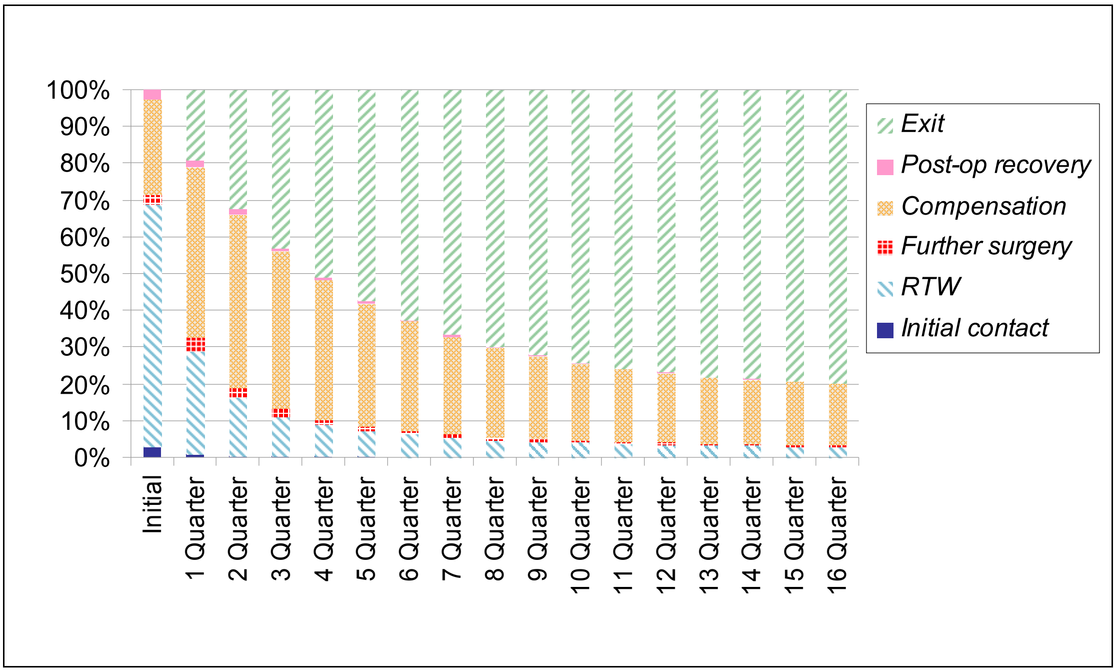

Figure 19: Breakdown of Agent B cluster distribution over time

Now that we have examined the scheme and two selected agents in detail, we again approach the question of what is an optimum claim management approach; and what might it look like. The approach we have applied assumes that 'someone' is already applying the best practice claims management, at any one point in the claims process. If we have captured the essence of that through the combination of the agents claim path and the average cluster costs, then we should be able to 'select' the best practice from any one agent and apply that practice to the equivalent cohort of claims from above. This should provide us with a view of 'how much better it can be', with sharing of best practice across the scheme.

Figure 20 and 21 following provide the projections of these selected optimum paths and cluster costs. Figure 20 implies that scheme RTW rates can be at the level of the best agents, in the order of 93%, which is an improvement over the current projected scheme average of 90%. It also shows that the long term cost per quarter can be $15m compared to the current projected scheme cost of $23m per quarter. The speed of the claim paths also makes a difference, and the present value of the optimum approach is in the order of $0.96b compared to the projected scheme present value of $1.15b, representing a present value claims costs reduction in the order of 15%.

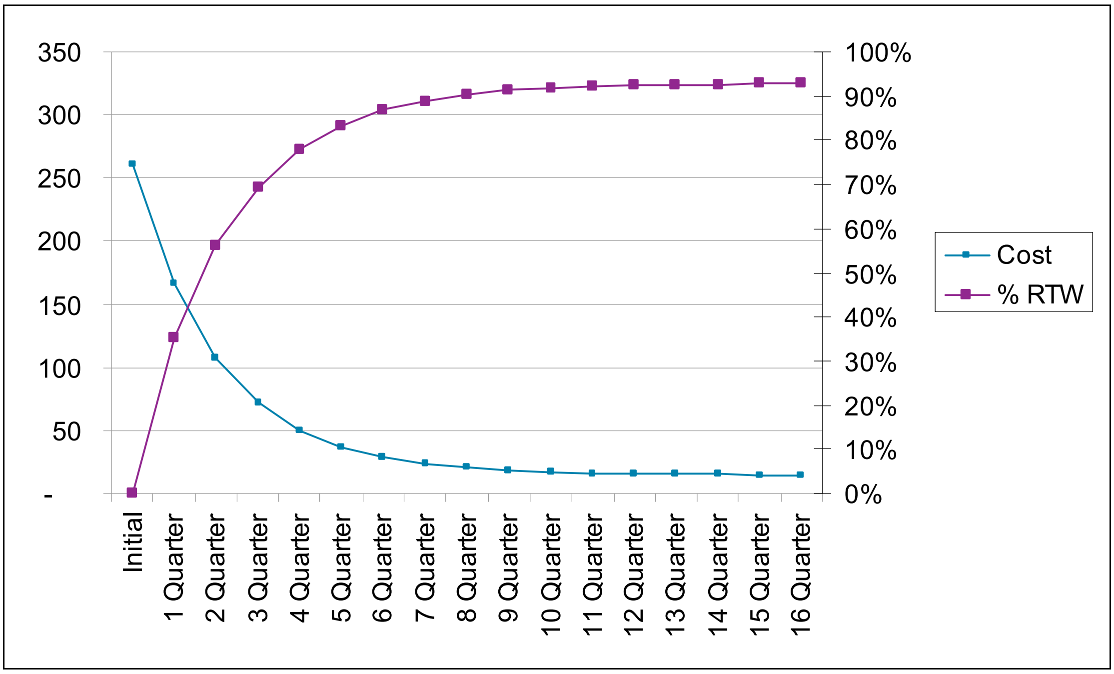

Figure 20: Optimum scheme RTW and total claim costs per quarter ($m)

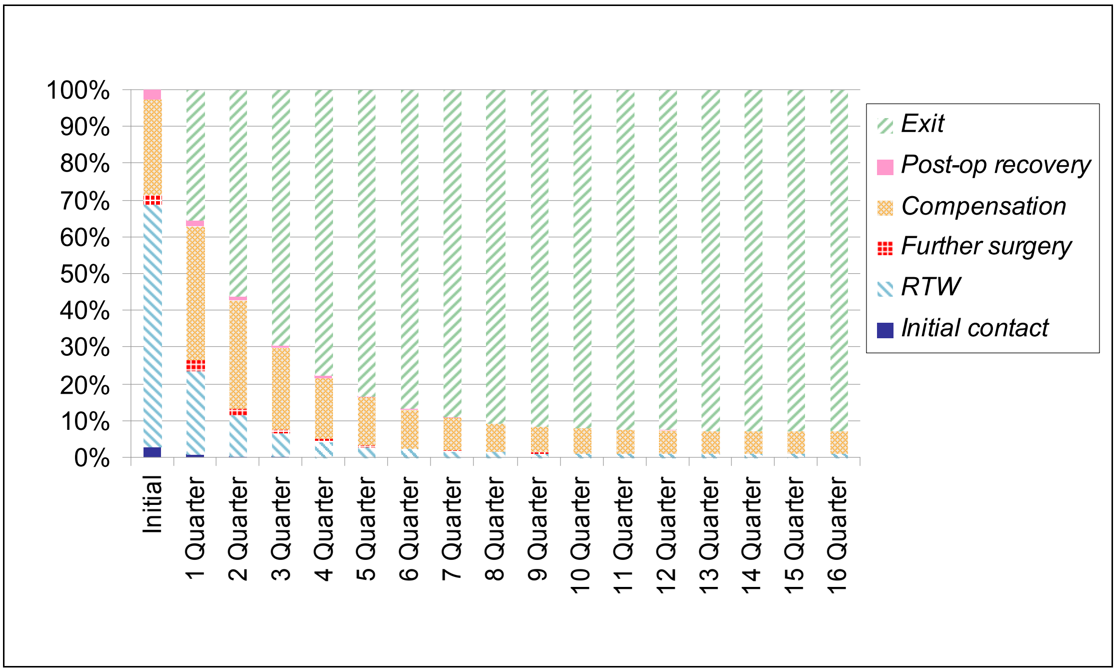

Figure 21: Breakdown of optimum scheme cluster distribution over time

Future steps

There are many aspects of the method, analysis and projection that could be improved to enhance the quality of the data directly as well as fine tuning the optimisation approach. We have listed several following:

- Adjusting the clusters to be sensitive to time since the injury occurred. This would imply that the transition rates would also vary by development period
- Identifying additional outcome measures for comparing claim management performance across the scheme
- Adding finer granularity in the clustering process (adding more clusters that have more specific treatment regimes embedded). This may not necessarily assist in claims strategy

development, but it would assist in identifying and fine tuning specific claims management processes

- Understanding why claims behaviour is exhibited by workers; i.e. why a claimant visited a GP and not the Public Hospital emergency department. It may be reasoned that most accidents are minor in the beginning and so the aim of this understanding would be to prevent escalation of claim severity due to poor claims management practices
- Investigating whether it is worthwhile prolonging certain claims management regimes to reduce the overall cost of the claims; and identifying which are the preferred areas to address to achieve the desired result
- Investigating if prolonged time in specific cluster alters the cost and later claims paths related to that claim i.e. the longer a claim stays on Weekly payments, the harder it may become to achieve a return to work outcome
- Adding the policy data to the input to the clustering to enable key attributes around industry and employer size to be isolated and examined
- Identifying key behaviours of groups of providers in the scheme and informing strategies to maximise the efficacy with which the scheme can work with these providers

Conclusions

The objective of this paper was to outline an analytical framework to understand the claims management processes within a scheme so that overall improvement strategies can be devised, while at the same time attempting to retain the granularity to be able to identify particular claims management tactics with which to implement the strategies.

We have found that the clustering technique provides a way of understanding claim patterns giving a good mix between the overview and the detail of the claims activities.

We have also found that the transition rates provide a way of understanding the differences between differing claims paths, whether that is comparing agents or different injury types.

Overall, the use of the best practice from any one agent could produce reduced life costs of claims in the order of 15%, assuming that the same types of claims are entering the scheme. The next step in the optimisation process is to identify the best approach across many different injury types, which may allow savings in excess of 15% across the scheme. The current best performing agent is likely to be improved further by this approach.

This framework, like most analytic tools, is designed to assist in targeting and taking the first steps on the path to improving claims management practices and a deeper understanding of the existing claims management methods needs to be first gained by examining the underlying behaviour within the claims management functions. This approach is well suited to performing the function of identifying where to look next.

Applications

We have identified the following applications for this framework:

- In development of treatment guidelines for medical practitioners, in particular GPs, that take into account the operation of the scheme within which the claimants are participating. For example, a GP may select a preferred treatment plan based on medical opinion, but given the constraints of the scheme, an alternate treatment plan may provide a better outcome for the claimant
- An approach to health care for use in the broader national and state health system, where better paths of care are identified, given the current state of the patient
- Informing the better approaches to claims management based on the current injury progression, for example, within a Tiered claims strategy where injury type, severity and location are key drivers of the adopted strategy
- Informing approaches where the average time to exit the scheme is minimised, i.e. fast exit strategies based on injury types and severities
- Using the framework as a way of comparing claims management practices of agents against a benchmark; which would be able to allow directly for differences between characteristics of employers and the mix of claims, so that performance was measured on a like for like basis
- Using the framework as a way of comparing internal claims management teams (within an agent); which would then allow for transferring the identified best skills or strategies across the remaining teams within an agent
- Applying a cost benefit approach to claims management to determine if additional claims expenditure will produce lower overall claims expenditure
- To provide greater insight into valuing scheme liabilities, that reflects better the actual claims process rather that applying a payment type model

Acknowledgements

We wish to acknowledge WorkCover NSW for providing access to the base data used in the development of this paper and Foo Xin from Ernst & Young for his technical support.