Chapter 5

Results and Discussion

In this chapter, we present results gathered from testing our implemented domain. For ease of reference, we begin by repeating our initial testing-based research sub-questions and the associated criteria for success (see Table 1.1, page 16). If the model satisfies these criteria for our implemented domain, then the model can be said to succeed in achieving the overall goal: to generate distinct personalities with a minimum of handcrafting. Further, the criteria span the different ways that personality can be measured according to Caspi & Roberts (see Section 2.3, page 66).

1. How does *adaptation* affect character behaviour?

    (a) Does behaviour change over time? *Behaviour* changes over time.
    (b) Can characters learn about specific, functional goals? When given a functional goal to learn, the majority of characters choose the "correct" action the majority of the time, based on *behaviour*.
    (c) How does reward change with time? *Reward* values are on average higher using our model than when random choice is used.
    (d) What happens if adaptation is turned off? Compared to when adaptation is turned off; both *individuality* and *reward* are higher when adaptation is used.

2. How does *context* affect character behaviour?

    (a) Does character behaviour differ in different contexts? For one character's *behaviour*, show that in different contexts the action chosen the majority of times is different.

(b) What happens if context is turned off? Compared to when context is turned off; both *individuality* and *reward* are higher when context is used.

3. How can personality be implemented so that the same template can be used to create a number of distinct, *individual* characters, according to their behaviour?

   (a) Are the behaviours of characters different from each other over time? Character *behaviour* passes the chi-squared test.
   (b) Are any individuals obtained? Based on their *individuality*, at least one character is different from the majority of the other characters when they are all based on the same template.

In this chapter, we address all of the above criteria for success in sections based on the three measures of effectiveness used to calculate the results. We start addressing the criteria for success by presenting results that relate to character behaviour (particularly while using the full model, 'normal' mode). This is followed by personal reward, and finally the quantitative individuality measure. We then present results relating to personal reward and individuality which were measured across all Cases and all modes. After presenting these results, we discuss implications of the results. The final section identifies some extra findings and several interesting results that emerged during testing, particularly in reference to the domain-dependent knowledge (friendships and happiness) used by characters. We finish with a summary of results.

5.1 Behaviour

We will now look at results from the 'normal' mode of each Case. We examine specific example runs to investigate in detail how the characters' actions change over time, how they learn about sub-plans based on their soft goals, how they behave in different contexts; whether they are different, and how individual characters compare to each other. That is, we are addressing the criteria for success (see Table 1.1, page 16) that relate to character *behaviour*.

5.1.1 Behaviour Over Time (Research Sub-question 1a)

In this section we are testing to determine whether behaviour changes over time. We found that for all Cases the characters' behaviour did change over time and responded

to their dynamic environment. We examined the number of times actions were chosen for intervals over the entire simulation in 'normal' mode runs for every case. We took a single sample run (out of the 10 possible runs) from each Case and graphed each character on a separate graph. The full set of graphs from these sample runs can be found in Appendix A (page 193). Here, we show two graphs for each Case from two different characters. For example, Figure 5.1 shows the graphs for Anna (Figure 5.1(a)) and Deb (Figure 5.1(b)). In these graphs, simulation time is on the x-axis. Each line represents a different top-level activity ("Insult", "Move" and "Wait"), and the y-axis shows how many times that particular plan was chosen over the data collection time period. That is, between each data output step, we summed the number of times the character choose each activity, and this is the value shown on the y-axis. The graphs illustrate how each character's most chosen activity changes over time.

Case 1: Clear Preference Against One Activity The graphs in Figure 5.1 are taken from a single randomly chosen run of Case 1 using the 'normal' mode. In this Case, the soft goal personalities all have a clear preference against one activity, insults. We can see in the sample graphs that, after an initial period of learning, both the characters chose "insult" the least frequently[1]. After this time, approximately half of the characters learned that "wait" is the most desirable activity, while the other half believe that "move" is the best activity. We can see that Anna chose "wait" most frequently (Figure 5.1(a)), while Deb chose "move" most frequently (Figure 5.1(b)).

Case 2: Multiple Ways to Achieve Goals Two character graphs from Case 2 are shown in Figure 5.2. Although some characters did not have a clearly preferred behaviour (such as Deb, Figure 5.2(a)), we also show here one character whose most chosen behaviour stabilised to be "insult" (Gina, Figure 5.2(b)). In this Case, since there were multiple ways to achieve the goals, some characters (such as Deb and others shown in Figure A.2, page 195) never found a most preferred activity. These characters changed their behaviour to match their environment, while other characters found a stable strategy that worked for them.

[1]This is true for all characters, as can be seen from the full set of graphs for all characters in the appendices, Figure A.1 (page 194)

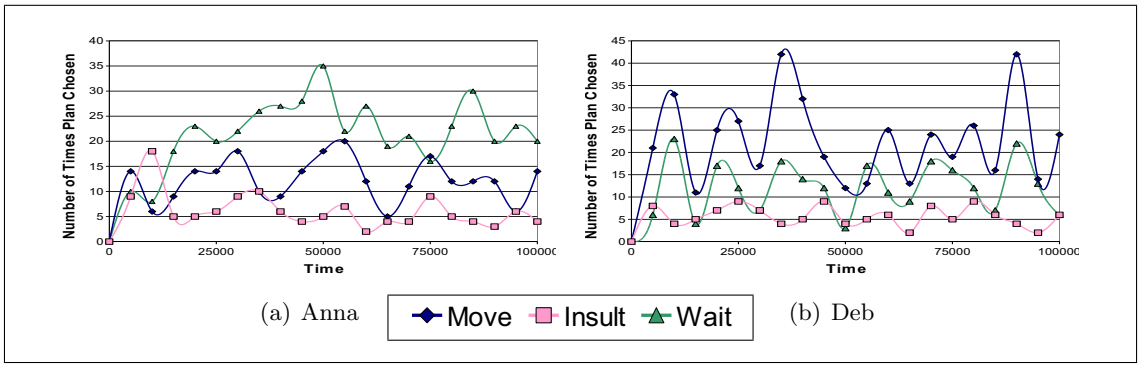

Figure 5.1: Sample character graphs for Case 1 of behaviour based on the individual: Action choices for each character for a particular run of Case 1 (Clear Preference Against One Activity) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Each line represents a different activity.

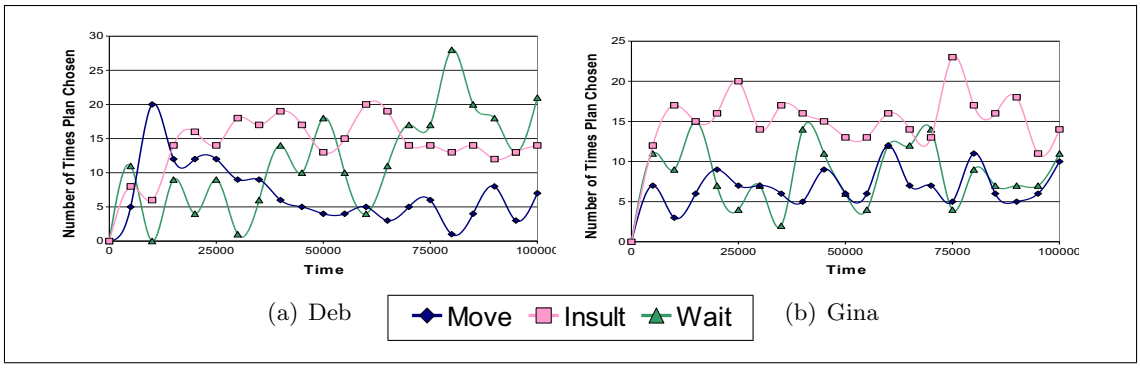

Figure 5.2: Sample character graphs for Case 2 of behaviour based on the individual: Action choices for each character for a particular run of Case 2 (Multiple Ways to Achieve Goals) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Each line represents a different activity.

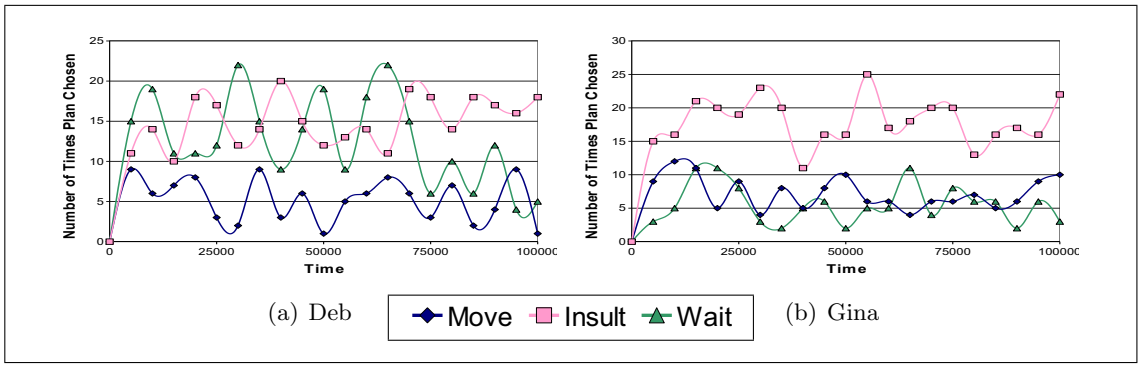

Figure 5.3: Sample character graphs for Case 3 of behaviour based on the individual: Action choices for each character for a particular run of Case 3 (Conflicting Goals) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Each line represents a different activity.

Case 3: Conflicting Goals Sample graphs from Case 3 are shown in Figure 5.3. In this Case, characters are trying to insult others, but not be insulted themselves. The full set of graphs for the characters in Figure A.3 (page 196), shows that Anna, Fran and Gina (also shown here in Figure 5.3(b)) all have a clear and early tendency towards "insult". All the other characters eventually settled on choosing "insult" most frequently, but they have some uncertainty over longer time periods. This can be seen for Deb in Figure 5.3(a), where she is alternating between "wait" and "insult".

Case 4: Complex Soft Goal Personality In Case 4, none of the characters show a clear tendency towards any of the three top-level activities, see Figure 5.4. A character's behaviour fluctuates while learning which plan will achieve their soft goals best for that particular context in a changing environment. Since character behaviour is dependent on what the other characters choose, it is possible that each individual character cannot work out the best strategy against the other characters because those characters are changing their strategies as well. That is, the character's fluctuations in behaviour could be due to the fluctuations of other characters.

Case 5: Different Soft Goal Personalities The behaviour graphs in Figure 5.5 show an example run from Case 5 where characters have different soft goal personalities.

- Anna and Bec (Figure 5.5(a)) have the same soft goal personality as Case 4, that is, they are trying to achieve six soft goals simultaneously. As we found for the

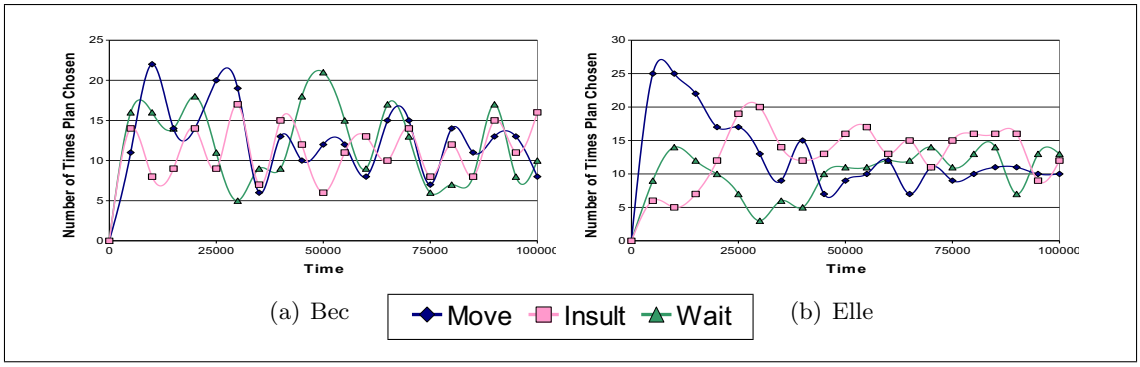

Figure 5.4: Sample character graphs for Case 4 of behaviour based on the individual: Action choices for each character for a particular run of Case 4 (Complex Soft Goal Personality) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Each line represents a different activity.

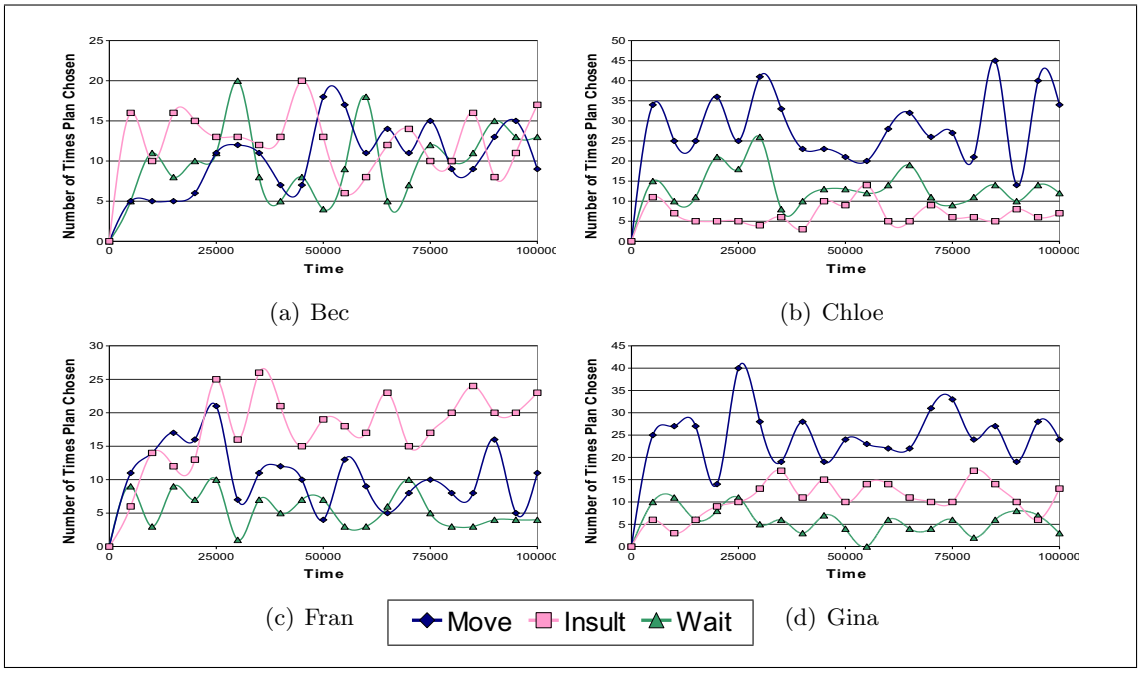

Figure 5.5: Sample character graphs for Case 5 of behaviour based on the individual: Action choices for each character for a particular run of Case 5 (Different Soft Goal Personalities) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis.

characters in Case 4, Bec's behaviour fluctuates and there is no clear tendency towards any activity.

- Chloe (Figure 5.5(b)) and Deb have the same soft goal personality as in Case 1, a clear preference against one activity, "insults". Similarly to the characters in Case 1, Chloe has a low tendency towards "insult", and has a clear tendency towards "move".

- Elle and Fran (Figure 5.5(c)) have the same soft goal personality as Case 3, conflicting soft goals. Fran chooses "insult" most frequently, which she must believe will allow her to insult others, while minimising being insulted herself.

- Gina (Figure 5.5(d)) and Heidi have only one soft goal: to make friends. Gina has a clear tendency towards "move" for the entire simulation. Heidi fluctuates most frequently between choosing "insult' and "move".

5.1.1.1 Comparison to Expectations

In Section 4.2.5.1 (page 138), we outlined our expected results in response to the research sub-question 1a relating to behaviour. We now consider how the model performed compared to these expectations:

1. Clear Preference Against One Activity: characters did learn not to choose "insult" as often as the other activities, as was expected.

2. Multiple Ways to Achieve Goals: some characters chose a particular activity more often than the other activities, as was expected. However, many characters did not have a clear tendency towards any action.

3. Conflicting Goals: several characters had a clear tendency towards "insult". The other characters alternated back and forth between "insult" and another activity, as was expected.

4. Complex Soft Goal Personality: the characters' plan choices changed so that the characters often chose each of the activities approximately the same number of times. This was not as expected and may mean that characters were unable to develop clear activity preferences.

5. Different Soft Goal Personalities: character behaviour did depend that particular character's soft goal personality. For example, the characters given soft goal preferences away from "insult", show a tendency not to choose that plan. This means that, as expected, not all characters had the same most chosen actions.

We expected that the characters would be able to adapt to their environment and change their behaviour. Characters did indeed change their behaviour based on the other characters in their environment.

5.1.2 Learning A Functional Soft Goal (Research Sub-question 1b)

We now examine whether the characters are able to learn more subtle preferences at the sub-activity level, rather than the activity level. In this section we address the research sub-question focussing on whether characters can learn about specific, functional goals. We consider in particular whether, when given a functional goal to learn, the majority of characters choose the "correct" action the majority of the time, based on *behaviour*. It should be remembered that in the current implementation, we used a bucket choice threshold so that the most desirable activity group was chosen 60% of the time (see emotionality values used in Section 4.2.2, page 122). This means that, even if a character places its highest preference on the correct behaviour, it is likely to choose this correct behaviour only 60% of the time. The bucket choice threshold was intended so that characters would make a suitable trade-off between exploiting a known successful action versus exploring other possible action.

This criterion relates specifically to Case 4, which allows the characters to learn about functional soft goals. In Case 4, the characters had six different soft goals they were pursuing. Two of these soft goals were: "be close to friends" and "don't be close to enemies". If the characters were achieving (or learning) these goals, then they should be choosing the actions "move towards friend' and "move away from enemy". That is, the characters should be able to learn that these two plans directly achieve each of those two soft goals respectively.

Within the goal/plan hierarchy, the characters cannot choose directly between "move towards friend" or "move away from friend". In the given goal/plan hierarchy (see Figure 4.2, page 106), the characters are able to choose first which direction they want to move and next who they want to move with respect to. For example,

if Deb has already chosen to "move towards" someone, she can now choose between "move towards a friend", "move towards an enemy" or "move towards a neutral". Although, if Deb currently has no enemies (for example), she will be unable to choose the plan "move towards an enemy", since the plan will not be applicable.

In Figure 5.6 we show the number of times an example character, Anna, chose each plan over the simulation with a graph for each of the categories of 'friend', 'neutral', or 'enemy'. Results from all characters for this particular run can be found in Appendix B (page 199). In the graphs each line represents the number of times that Anna chose move towards or move away.

In Figure 5.6(a), we can see that the line representing "move towards friend" is chosen more times than "move away from friend". This means that Anna has learned how to achieve the soft goal "be close to friends". This was verified in results that showed that Anna chose "move towards friend" more frequently compared to "move away from friend". In Figure 5.6(b), Anna appears to have no clear tendency to move towards or away from characters. This is most likely due to the fact that no soft goals relate to these plans, so the character can choose any plan and it will not affect its achievement of soft goals. Figure 5.6(c) shows the plans that relate to moving with respect to an 'enemy'. In this situation, Anna appears to prefer to "move towards enemy", which is the opposite of what she should be learning. This means that Anna did not learn about the soft goal "don't be close to enemies".

We now examine the results in relation to the specific criterion. For characters to learn this soft goal, the majority of characters should choose the "correct" action the majority of the time (i.e. over 50% of the time). Here, there are two correct actions: "move towards friend" and "move away from enemy". Across the 10 runs of this Case in 'normal' mode, an average of 6.1 ± 1.1 characters chose "move towards friend" the majority of the time (compared to "move away from friend"). In some instances, a few characters did not ever have friends, so they were never able to choose to move with respect to a friend. An average of 5.9 ± 1.2 characters chose "move away from enemy" the majority of the time (compared to "move towards enemy").

5.1.2.1 Comparison to Expectations

It was expected that the results would satisfy this criterion. To reach a majority, the average number of characters must be greater than 4 characters. Based on the

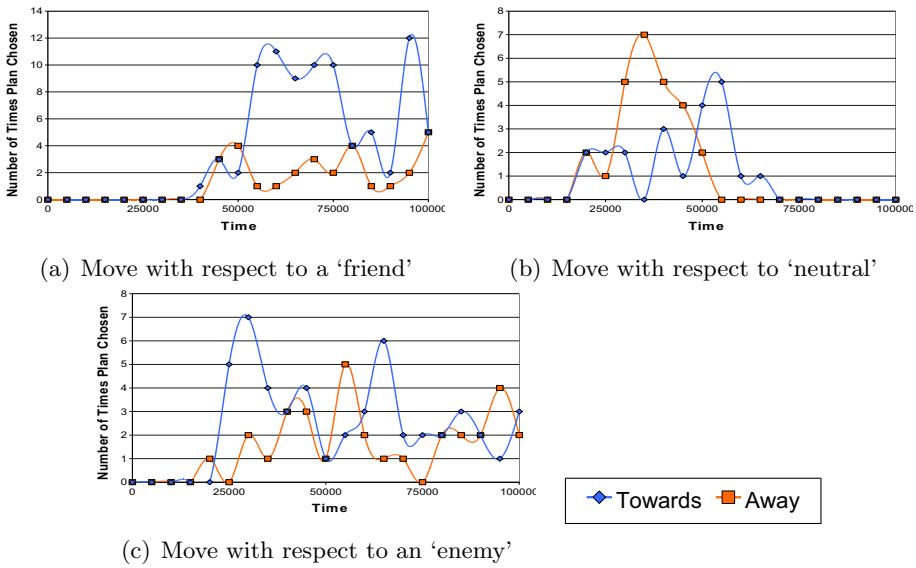

(a) Move with respect to a 'friend'

(b) Move with respect to 'neutral'

(c) Move with respect to an 'enemy'

Figure 5.6: Anna's Behaviour when moving for Case 4: Action choices for Anna for a particular run of Case 4 (Complex Soft Goal Personality) 'normal' mode. In each graph, the number of times the character chose each of "move towards" or "move away" over the time interval is shown on the y-axis.

behaviour data presented above, for both possible "correct" actions, the majority of characters *did* choose the correct action the majority of the time.

When we examined character behaviour over time in Figure 5.4 (page 150), we noticed that the characters did not have clear tendencies towards any of the top-level activities. Despite this apparent lack of learning at the top-level, reaching the criterion has established that the characters can learn appropriately about the lower-level plans based on their specific soft goals.

5.1.3 Behaviour in Different Contexts (Research Sub-question 2a)

This section answers the research sub-question: Does character behaviour differ in different contexts? The criterion for this question is: for one character's *behaviour*, show that in different contexts the action chosen the majority of times is different.

To test this, we need to consider each of the Cases and determine whether any of the characters have different behaviour. For each Case and every character from the sample run, we looked at the actions (in particular, activities) chosen and the context the character was in when they chose that activity. We used data counts of the number

times each particular character was in each context. Based on this data, we choose the top two contexts to examine in detail. For these contexts, we examined whether one of the three top-level activities was chosen the majority of the time. That is, when the character was in a particular context (e.g. "HH"), we examined whether any of the activities ("move", "insult", "wait") were chosen more than 50% of the time. If an activity was chosen the majority of the time, how did this activity compare to the choices in the second most frequently occurring context. If, in the second context, there was no majority for any activity, or the majority activity was different; then it could be said that the character's behaviour was different in those two contexts. That is, the criterion is met for that Case. It was unnecessary to test Case 5, since this Case includes characteristics that are can be found in the other Cases.

We found that, in all of the sample runs considered, there were a few characters whose behaviour was different in different contexts, i.e the criterion to obtain at least one character was met. To demonstrate these results, we will present sample graphs for four example characters from each of the four sample runs for the first four Cases. For each character, we show two graphs of behaviour in each of its two most used contexts, for example Figure 5.7. These graphs will indicate that for that particular character, the action chosen the majority of times is different.

Case 1: Clear Preference Against One Activity From Table 4.2 (page 136), the soft goals the characters are trying to achieve are:

1. "Don't be insulted"; and
2. "Don't insult people".

This means that, their ideal context is when they are not being insulted, and are not insulting others. In our domain, this is represented by the context: "HL". Figure 5.7 shows Bec in her two most common contexts: "HH" (Figure 5.7(a)) and the ideal context, "HL" (Figure 5.7(b)). When Bec is in the ideal context ("HL") we can see that the activity that she chooses the majority of the time is "wait", see Figure 5.7(b). In Bec's other most visited context (when she is insulting many people), her favourite activity is "move", see Figure 5.7(a). We note, that in both contexts, Bec chooses "insult" the least often. This is in line with her soft goal personality that ensures she should have a preference against insults.

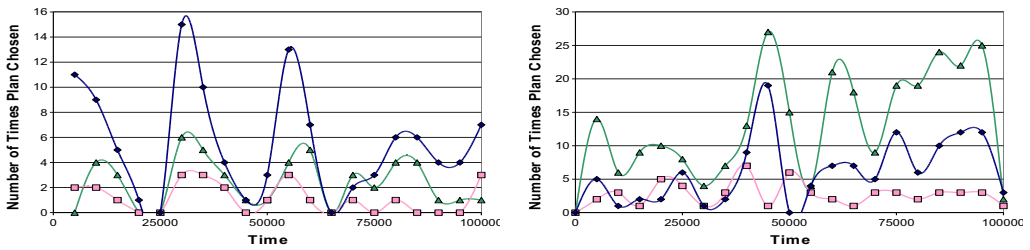

(a) Context HH: Bec is not being insulted and is insulting people a lot

(b) Context HL: Bec is not being insulted, but is hardly insulting anyone herself (ideal context)

Figure 5.7: Sample character behaviour graphs for two contexts in Case 1: Action choices for Bec for two contexts for a particular run of Case 1 (Clear Preference Against One Activity) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Note for this Case the ideal context is: "HL".

Case 2: Multiple Ways to Achieve Goals From Table 4.2 (page 136), the soft goals the characters are trying to achieve are:

1. "Don't be insulted"; and
2. "Make friends".

This means that, in our implementation, the ideal context is "HH", which represents when the character is not being insulted and has many friends. The two contexts that Chloe was in most frequently were: "HM" and "LM". The difference between these contexts is how frequently Chloe is being insulted. In both contexts, she has a moderate number of friends. In "HM", she is not being insulted much. The context "HM" represents the closest Chloe comes to achieving her soft goals most of the time. The graphs in Figure 5.8 show that, in the context "HM", Chloe chooses "wait" a clear majority of the time (Figure 5.8(a)). In the context "LM", there is no clear majority. Chloe's behaviour is different because, in one context, she clearly chooses one plan above others and in the other context she fluctuates with no clear preference.

Case 3: Conflicting Goals From Table 4.2 (page 136), the soft goals the characters are trying to achieve are:

1. "Don't be insulted"; and

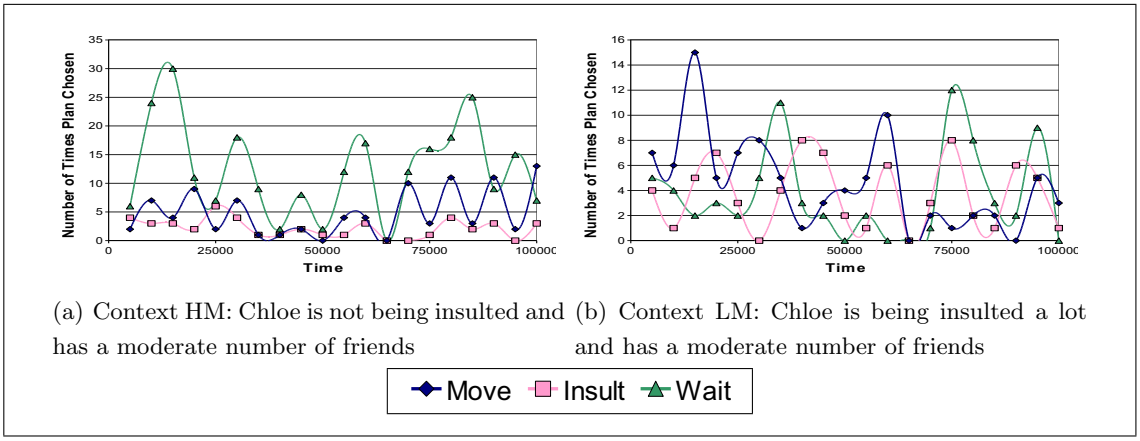

(a) Context HM: Chloe is not being insulted and has a moderate number of friends

(b) Context LM: Chloe is being insulted a lot and has a moderate number of friends

Figure 5.8: Sample character behaviour graphs for two contexts in Case 2: Action choices for Chloe for two contexts for a particular run of Case 2 (Clear Preference Against One Activity) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Note for this Case the ideal context is: "HH".

2. "Insult people".

For this Case, the ideal context is "HH", which Deb is in very frequently. Her other most frequent context is "LH", which represents the Case where Deb is being insulted frequently and is insulting others a lot. When Deb is not being insulted frequently ("HH"), she chooses "wait" the majority of the time, see Figure 5.9(a). When Deb is being insulted frequently ("LH"), she chooses "insult" most frequently. That is, Deb has learnt that, when she is being insulted, her best activity is to insult others (perhaps she has found that this stops them insulting her).

Case 4: Complex Soft Goal Personality From Table 4.2 (page 136), the soft goals the characters are trying to achieve are:

1. "Don't be close to enemies";
2. "Be close to friends ";
3. "Don't be insulted";
4. "Insult enemies";
5. "Don't make enemies"; and
6. "Make friends".

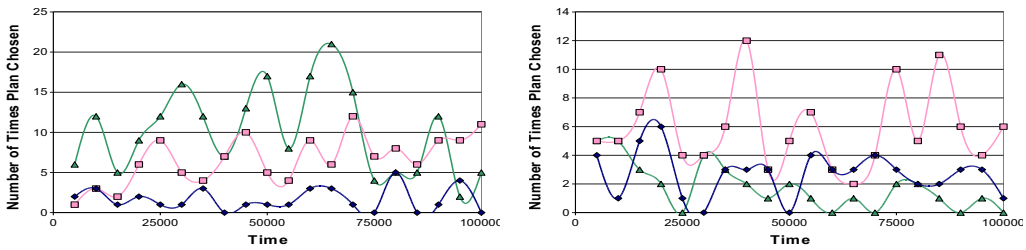

(a) Context HH: Deb is not being insulted and is insulting people a lot (ideal context)

(b) Context LH: Deb is being insulted a lot and is insulting people a lot

Figure 5.9: Sample character behaviour graphs for two contexts in Case 3: Action choices for Deb for two contexts for a particular run of Case 3 (Clear Preference Against One Activity) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Note for this Case the ideal context is: "HH".

The alphabetical ordering of these contexts is important because it allows us to interpret the context strings. The ideal context in this Case is: "LHHHLH". The two contexts that Gina was in most frequently were: "LLHLHM" and "LLLLHM". These two contexts differ in the amount that Gina is being insulted. In both contexts, Gina is neither close to her enemies nor her friends, is hardly insulting her enemies, has many enemies and a moderate number of friends. In "LLHLHM", Gina is not being insulted (more ideal), whereas in "LLLLHM", Gina is being insulted a lot (not ideal). If we examine Figure 5.10, we can notice that for Gina is rarely choosing any activity in either context. When we examined which context Gina was in most, we found that Gina's context changed so frequently that she hardly had two clear contexts she was in most. When we ran the simulation for longer, we found similar results, that is, the characters continued to change contexts frequently. This is a problem that is addressed in the discussion of results that follows. From the graphs presented here, we can notice that, similarly to Deb in Case 3, Gina learns that, when she is being insulted a lot, she should insult others, see Figure 5.10(b). This behaviour for this context is different from Gina's behaviour for her other context where she has no clear preference between the activities but perhaps has a slight tendency towards "move".

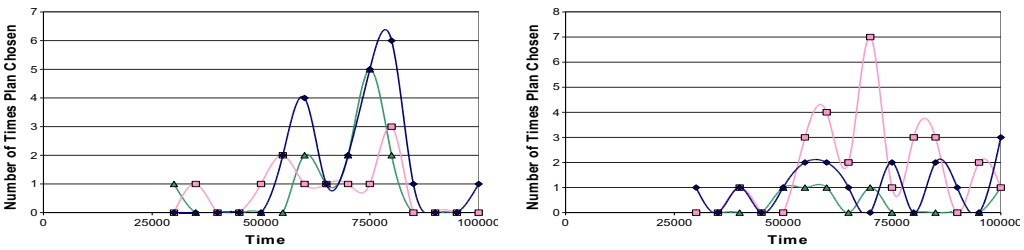

(a) Context LLHLHM: Gina is not close to her enemies nor her friends; she is not being insulted; is hardly insulting her enemies; has many enemies; and a moderate number of friends

(b) Context LLLLHM: Gina is not close to her enemies nor her friends; she is being insulted a lot; is hardly insulting her enemies; has many enemies; and a moderate number of friends

Figure 5.10: Sample character behaviour graphs for two contexts in Case 4: Action choices for Gina for two contexts for a particular run of Case 4 (Clear Preference Against One Activity) mode 'normal'. In each graph, the number of times the character chose each of the three top-level activities is shown on the y-axis. Note for this Case the ideal context is: "LHHHLH". Also note that the difference between these contexts is the degree to which Gina is being insulted.

5.1.3.1 Comparison to Expectations

We now compare the results to our expectations that were outlined in Section 4.2.5.1 (page 139).

1. Clear Preference Against One Activity: we believed this Case would not satisfy the criterion since there were less ways a character could differ. However, we did find some characters whose behaviour is different in different contexts, so the criterion is reached.

2. Multiple Ways to Achieve Goals: we expected this Case to reach the criterion and it did.

3. Conflicting Goals: as expected, this Case satisfied the test.

4. Complex Soft Goal Personality: we were uncertain whether this Case would succeed, but our results for Gina show that it did.

5. Different Soft Goal Personalities: since the other four Cases succeeded, this Case will satisfy the criterion as well.

Research Sub-question	Case				
	1	2	3	4	5
1a: Behaviour Over Time	Success	Success	Success	Success	Success
1b: Learning A Functional Soft Goal	N/A	N/A	N/A	Success	N/A
2a: Behaviour in Different Contexts	Success	Success	Success	Success	Success
3a: Chi-squared Test	Success	Success	Success	Success	Success

Table 5.1: Results of testing Cases for behaviour-based criteria for success: Success or fail relates to results satisfying the criteria as detailed in Table 1.1 (page 16).

5.1.4 Chi-squared Test (Research Sub-question 3a)

The sub-question relates to whether behaviour of characters differs to other characters. The criterion was that character behaviour passes the chi-squared test. The chi-squared test is a standard statistical tool that establishes whether the behaviour of one character can be used to predict behaviour of another character. If the characters pass the chi-squared test, then the p-value confidence in the result should be < 0.05. This criterion relates to the 'normal' mode, so we applied the chi-squared test to characters in every run and every Case and found that the test was passed every time.

5.1.4.1 Comparison to Expectations

Our expectations for this criterion (see Section 4.2.5.1, page 140) were that Cases 2-5 would succeed, but Case 1 was only 'likely' to satisfy the test. This means our expectations were exceeded because all five Cases satisfied the test.

5.1.5 Summary of Behaviour Results

The results when using *behaviour* to asses the performance of characters in the model are summarised in Table 5.1. Comparing this to the expected results, Table 4.3 (page 138), it is seen that our model performed better than expected on all behaviour related criteria for success, since no Case failed any test.

5.2 Personal Reward

Every time a character completed an evaluation step, it output its personal reward value at that time. By examining the graphs of reward versus time, we found large

Case	Mode			Average
	Adaptation Off	Context Off	Normal	
Case 1: Clear Preference Against One Activity	-0.11	0.55	0.4	0.28
Case 2: Multiple Ways to Achieve Goals	-0.03	0.08	0.09	0.05
Case 3: Conflicting Goals	0.29	0.31	0.33	0.31
Case 4: Complex Soft Goal Personalities	0.01	0.06	0.05	0.04
Case 5: Different Soft Goal Personalities	0.01	0.23	0.19	0.14

Table 5.2: Average reward values for test Cases: personal reward scale is [-1,+1].

Case	Mode		
	Adaptation Off vs Context Off	Context Off vs Normal	Adaptation Off vs Normal
Case 1: Clear Preference Against One Activity	Yes	Yes	Yes
Case 2: Multiple Ways to Achieve Goals	Yes	Yes	Yes
Case 3: Conflicting Goals	No	No	Yes
Case 4: Complex Soft Goal Personality	Yes	No	Yes
Case 5: Different Soft Goal Personalities	Yes	Yes	Yes

Table 5.3: Significant differences between personal reward values for different test modes.

fluctuations in personal reward values. Due to this, we used an average of the personal reward as the test datum and not the final reward obtained. We collated the reward data across all the Cases and modes to generate the results shown in Table 5.2. Each value in the table represents the average of 80 data points (based on 8 characters in 10 runs), where each data point is the average for each character of approximately 1000 reward calculations over the simulation run. Given the large amount of data, normality tests were not considered necessary.

Since these values are averages, we examined whether these values were statistically significantly different ($p < 0.05$) from each other. We found that in the 'normal' mode reward average was always significantly higher than in the 'adaptation off' mode. The results for statistical significance comparing the modes to each other are in shown Table 5.3. The relevance of this significance data is discussed in Section 5.2.2 (page 162).

We now discuss these results in relation to the criteria for success relating to reward

values. We begin by comparing reward values to random choice and then examine the effect of adaptation and context on reward values.

5.2.1 Reward Compared to Random Choice (Research Sub-question 1c)

The criterion for addressing this sub-question is that reward values are on average higher using our model than when random choice is used. Table 5.2 shows that the average reward values from testing across all Cases and all modes. For the mode using our model, 'normal' mode, the average reward values are greater than those generated using random choice, 'adaptation off' mode, in all Cases tested. Further, the differences were significant in all Cases, meaning that the model has satisfied this test in all Cases.

5.2.1.1 Comparison to Expectations

Our expectations (see Section 4.2.5.2, page 140) were that our model would generate higher average reward values than random choice. So our expectations have been confirmed and the model satisfies this test.

5.2.2 Effect of Adaptation and Context on Reward (Research Sub-questions 1d and 2b)

The criteria for 1d and 2b (as specified from Section 1.1, page 16) are:

- 1d: Compared to when adaptation is turned off, reward is higher when adaptation is used.
- 2b: Compared to when context is turned off; reward is higher when context is used.

We begin by considering and comparing the reward values across the modes and then across Cases. After this we summarise the results from testing in relation to the criteria and compare to what we expected to find.

5.2.2.1 Reward Across Modes

Effect of Adaptation on Average Reward Based on raw averages in all Cases, the 'adaptation off' mode, performed worse than the other two modes. The difference was significant for all Cases except Case 3, where difference between 'adaptation off' and

'context off' modes was not significant. This means that, in Case 3 (conflicting goals), without the assistance of context, characters would have been better using random choice to make decisions, rather than learn which choice to make. In Case 3, the characters are trying to achieve two goals that are unlikely to be achieved at the same time, "insult people" and "don't be insulted" and are actually in conflict with each other. In this Case, context is needed to enable the character to learn that, when one goal is not being achieved, it should attempt to achieve the other soft goal. Although the average for 'normal' was higher, it was not significantly better than 'context off' mode. The average reward value for 'normal' was significantly higher than that in 'adaptation off' mode.

Effect of Context on Average Reward The 'normal' mode produced a higher average reward than 'context off' in two Cases as shown in Table 5.2, although this difference was only significant in one Case. In the other three Cases, the 'context off' mode produced a higher average reward than 'normal', although in one Case the difference was not significant.

5.2.2.2 Reward Across Cases

Case 1 showed the highest reward values for the two modes 'context off' and 'normal'. This was the Case where there was a clear preference against one activity. This meant that, at the top-level of choosing between activities, the characters had soft goals that directly indicated that one of the activities should not be chosen. Presumably, that meant it was easier for characters to learn which activities and plans were not good and therefore the characters got closer to achieving their soft goals, i.e. not giving or receiving insults.

Cases 2 and 4 with multiple ways to achieve goals and complex soft goals, respectively, showed the lowest reward values. This is probably because characters found it difficult to find a good reward path. With complex soft goals (Case 4), the characters are trying to achieve too many goals at the same time and are unable to determine which choice to make. The problem is hindered by the fact that reward is not guaranteed. So, although choosing an action like "move towards friend" should improve soft goals, it may not always work. While the character is moving towards someone, that someone could move as well and so the character may or may not get closer to their

friend, thus changing the reward value so that it may not reflect the intended action. It was hoped that over time they would learn, apparently the environment fluctuated too much for this to happen. The implementation was run for a reasonably long period of time and the preferences of the characters were not changing greatly towards the end of the period, even if their rewards were fluctuating greatly.

5.2.2.3 Summary of Personal Reward Results

Across modes, actual reward values were, ranked from highest to lowest reward value: 'context off'; 'normal'; and 'adaptation off'. In relation to research sub-question 1d comparing the 'normal' mode with 'adaptation off' mode, the model satisfied the criteria in all Cases. Reward was higher in 'normal' mode than when using adaptation. Considering sub-question 2b, comparing the average reward in 'normal' mode with 'context off' mode, reward was higher in one Case, equal in two Cases; and less in two Cases.

Across Cases, actual average reward values were, ranked from highest to lowest reward value:

1. Case 3 (conflicting goals);
2. Case 1 (clear preference against one activity);
3. Case 5 (different soft goal personalities).
4. Case 2 (multiple ways to achieve goals); and
5. Case 4 (complex soft goal personality).

The average value for Case 1 'adaptation off' is very low compared to 'context off' and 'normal' modes. This brings the average across modes for Case 1 down so much that it does not produce the highest average reward value.

5.2.2.4 Comparison to Expectations

Across modes it was expected that the order would be (highest to lowest): 'normal'; 'context off'; 'adaptation off' (see Section 4.2.5.2, page 141). Actual results showed the average reward for 'context off' as being higher. This could be because context does not actually help the characters learn as effectively as was hoped.

Across Cases, we expected Case 1 to obtain the highest reward and Cases 3 and 5 to obtain the lowest reward values. Contrary to expectations, Cases 2 and 4 had

the lowest reward values. This may be because the characters were not able to find an optimal path from their many choices. The characters may have spent too long exploring their domain and the contexts available to them without exploiting paths that looked promising.

It was expected that Cases 3 and 5 would have the lowest reward values but this was not found to be the result. These are the Cases where there are multiple soft goal personalities in the model (Case 5) and where the characters have conflicting goals (Case 3). In the actual results, these Cases had moderate to high levels of reward. In Case 5, the higher than expected values could be because the individual characters found their own high personal reward values and were not as reliant upon others as expected. For Case 3 with conflicting goals, it appears that the characters were able to find a way to achieve their goals relatively well.

5.3 Individuality

The quantitative measure for individuality used was introduced in Section 4.2.3.3 (page 126). Individuality is determined from pair-wise comparisons of the behaviour of two characters at the same time intervals using paired t-tests. To ensure valid derivation of individuality, the data (action counts over time intervals) must first be checked to confirm that it is normally distributed. All the data we obtained was normally distributed.

A chi-squared test is also required to provide a preliminary check of whether there exists any difference between the characters. If a specific run does not pass the chi-squared test, there are no discernable differences between characters. Individuality gives a quantitative value to just *how* different the characters are, assuming that some difference has already been detected between characters. None of the Cases in 'adaptation off' mode passed the chi-squared test. This means that the behaviour in this mode was not unique to a particular character. This makes sense because, in the 'adaption off' modes characters were essentially using random choice to make decisions. This would be expected to result in non-discernable differences between the characters over time. Both 'context off' and 'normal' modes passed the chi-squared test in all Cases, allowing us to compare their results for individuality.

Case	Mode			Average
	Context Off	Normal	Statistically Different?	
Case 1: Clear Preference Against One Activity	20.0 ± 4.2	21.1 ± 4.3	No	20.55
Case 2: Multiple Ways to Achieve Goals	41.7 ± 3.4	36.8 ± 6.1	Yes	39.25
Case 3: Conflicting Goals	19.3 ± 5.6	26.5 ± 4.4	Yes	22.90
Case 4: Complex Soft Goal Personality	36.4 ± 4.6	14.1 ± 5.3	Yes	25.25
Case 5: Different Soft Goal Personalities	42.6 ± 3.4	38.5 ± 5.5	Yes	40.55

Table 5.4: Individuality for test Cases: average number of significant differences between individuals for test Cases. Maximum number of differences is 56.

Individuality uses paired t-tests to determine the number of significant differences between the actual action choices that the characters made for the three top-level activities: "move", "wait", "insult". We examined the percentage of times that they choose the plans over an interval since this is more observable to a user than the actual number of times (see Section 4.2.3.3, page 126). The maximum number of significant differences possible is $2 \times C_2^n$ (see Section 4.2.3.3, page 132). Since we are using eight characters, $n = 8$, we can find that the maximum number of difference between our characters is: $2 \times C_2^8 = 56$. If the individuality value obtained for a particular run was 56, it would mean that all eight characters made significantly different choices to every other character for all three top-level activities. Keeping this maximum value in mind we present the results across Cases and modes in Table 5.4, including whether the difference in each Case between the two individuality values are significantly different. These results show that we found the differences between the two modes to be significantly different in all but Case 1 (clear preference against one activity).

We now consider the criteria for success relating to individuality. We begin by considering the effect of adaptation and context on individuality. Then we consider how well our model performs in generating individual characters with behaviour that is very different from others, research sub-question 3b.

5.3.1 Effect of Adaptation and Context on Individuality (Research Sub-questions 1d and 2b)

The criteria for research sub-questions 1d and 2b (as specified from Table 1.1, page 16) are:

- 1d: Compared to when adaptation is turned off, individuality is higher when adaptation is used.
- 2b: Compared to when context is turned off; individuality is higher when context is used.

The 'adaptation off' mode did not pass the chi-squared test and therefore we can conclude that individuality is higher when adaptation is used. We now consider the effect of context on individuality and then how individuality varies across Cases. After this we summarise our results for the criteria and compare them to our expected results.

5.3.1.1 Individuality Across Modes

In this section, we consider the effect of context on individuality, i.e. testing research sub-question 2b, using Table 5.4. In Case 1 (clear preference against one activity), neither mode was clearly better or worse. This could be because the characters all learnt not to choose insults in both Cases, as indicated by the high reward values (indeed these two Case-mode combinations had the highest reward). Now, if they have learnt as they were supposed to, then they would eliminate one of the ways they can differ from each other. This means that perhaps the maximum individuality value is limited by the soft goals they are trying to achieve. In order words, the soft goals for this Case effectively reduced the number of ways that characters could demonstrate different behaviour. So, it is possible that both modes reached the limit and therefore are approximately the same.

In Cases 2, 4 and 5, 'context off' mode showed significantly more differences between characters than 'normal', i.e had better individuality. In Case 2 (multiple ways to achieve the goals), it appears that context did not help the characters choose their own different way to achieve the soft goals. We were unable to determine why this would be so for this Case compared to others. In Cases 4 and 5, some of the characters had a large number of contexts to learn about (all characters in Case 4 and some in Case 5). So perhaps they were not able to learn their preferred path effectively with the extra

contexts to learn about. Certainly, the reward values for Case 4 were very low (see Table 5.2, page 161). This inability to learn could be due to fluctuations in calculations of personal reward. This in turn could mean the characters ended up with no clear preference or behaviour and so could not easily be distinguished from each other.

It is interesting to note that Case 4 had a large difference in individuality between the modes (but curiously not between reward values). That is, when characters used contexts to make decisions, they found it exceedingly difficult to learn a clear preferred path, and so characters may have switched between activities (see for example the characters in Figure 5.4, page 150). On the other hand, with adaptation and no contexts, the characters could learn which plan to choose irrespective of the soft goal context, and so learnt their own way to differ from each other.

The individuality for the 'normal' mode was significantly higher than in 'context off' mode for one Case, where the characters had conflicting goals (Case 3). This indicates that the individuality was greater when using our full model with contexts. In Case 3, characters had conflicting soft goals, so that using context may help the characters work out which soft goal to achieve next, e.g. the one that is failing currently. Using context appears to have enabled the characters to build up their own unique way to achieving soft goals depending on the context. It shows that, in at least one Case, context-aware characters are more individual.

5.3.1.2 Individuality Across Cases

When comparing how the individuality varied across the different Cases, we use the average across the 'context off' and 'normal' modes. Based on this, we can see in Table 5.4 (page 166) that Case 5 generated the highest individuality. In Case 5, there were four starting soft goal personality templates used. This meant that from the beginning of the scenario the characters already had some differences between each other. So it makes sense that this Case produced the highest individuality. In fact, it would have been more surprising if this Case had not achieved the highest individuality.

Case 2, where characters had many different ways to achieve their soft goals, had the next highest individuality. In this Case, since their were so many possible optimal paths, characters found their own solution path that was different from the other characters, therefore providing greater potential for individuality.

The least significant differences are in Case 1, clear preference against one activity, and Case 3, where the characters had conflicting goals. It was expected that Case 1 would have a small number of differences between characters because the number of optimal plan paths available to the characters was diminished, as one activity (insults) is never desirable. Case 3, conflicting soft goals, had relatively high personal reward values (see Figure 5.2, page 161), so characters were achieving their goals. The low individuality value may be due to characters achieving their goals in the same way as the other characters.

5.3.1.3 Summary

Across modes, actual individuality values were, ranked in order of highest to lowest individuality: 'context off'; 'normal'; and 'adaptation off'.

Across Cases, individuality ranking was from highest to lowest:

1. Case 5 (different soft goal personalities);
2. Case 2 (multiple ways to achieve goals);
3. Case 4 (complex soft goal personality);
4. Case 3 (conflicting goals); and
5. Case 1 (clear preference against one activity);

5.3.1.4 Comparison to Expectations

Across *modes* we expected the 'normal' would have the highest individuality values since characters would have more dimensions in which to differ to other characters. However, it appears when characters could learn about contexts, they were less likely to have stable behaviour and therefore less likely to clearly differ from each other according to the individuality value. Across *Cases*, the actual ranking of individuality was exactly as expected (see Section 4.2.5.3, page 142).

5.3.2 Individuality per Character (Research Sub-question 3b)

The criterion for this test is: based on their individuality, at least one character is different from the majority of the other characters when they are all based on the same template. That is, we are considering each character's personal individuality measure, rather than the entire run's individuality measure.

Here, we are investigating individual characters who begin with the same template, i.e. Cases 1-4. In the previous section, we established that the characters are different from each other. However, if the characters are all only slightly different from each other, then perhaps the differences will not be obvious to an observer. If at least one character is significantly different from the majority of the other characters, then it is much more likely that there is at least one observable personality. For example, if the individuality value is 14, then there are 14 differences between the characters. This could mean that most characters are different from one of the characters on two dimensions (ie each character has almost 2 significant differences). On the other hand, it could mean that one character is different from every other character, and therefore that particular character has 14 differences. The second type of differences are probably easier for an observer or player to notice.

In order to consider the individuality value for each character, we separated out the paired t-tests for each character. This allowed us to count for each particular character how many differences existed between it and each of the other characters. The maximum is 14 differences, two for each of the 7 other characters. Remember that characters can only differ from each other on two dimensions, since we use the percentage of times each of the three top-level activities can be chosen (see Section 4.2.3.3, page 131). So, for a character to be different from the majority of the other characters, then that character would have more than 7 differences to the other characters. We then count the number of characters whose behaviour is different from the majority, giving us a value out of 8 (since there are eight characters). Table 5.5 shows the results for the average (across runs) of the number of characters whose behaviour is different from that of the majority of the other characters for both modes. The criterion states that at least one character should be different from the majority of other characters. So, to satisfy the criterion, the average number of characters should be greater than or equal to 1.

5.3.2.1 Comparison to Expectations

Our expectations for our model (i.e. 'normal' mode) (see Section 4.2.5.3, page 143) were that Case 1 was likely to fail and Cases 2 to 4 were likely to succeed. As seen in Table 5.5, Case 1 did fail the test, since the average was less than 1 character (0.7 ± 0.8), although some specific runs did satisfy the criteria. Cases 2 and 3 in 'normal' mode

Case	Context Off	Normal	Similar Value
Case 1: Clear Preference Against One Activity	0.7 ± 1.2	0.7 ± 0.8	Yes
Case 2: Multiple Ways to Achieve Goals	7.7 ± 0.9	6.0 ± 2.2	Yes
Case 3: Conflicting Goals	1.0 ± 0.9	2.4 ± 1.4	Yes
Case 4: Complex Soft Goal Personality	6.4 ± 1.8	0.5 ± 0.5	No

Table 5.5: Average number of characters who are different from the majority: maximum number is 8 meaning that all eight characters are different from the majority of other characters.

passed the criterion with values of 6.0 ± 2.2 and 2.4 ± 1.4 respectively. Case 4 did not succeed for in 'normal' mode (0.5 ± 0.5). However, for the 'context off' mode Case 4 did succeed the test well (6.4 ± 1.8). This could be due to the characters in Case 4 'normal' never having clear tendencies towards certain activities.

5.4 Discussion

Based on the results presented, it is possible to draw some conclusions about our theoretical model and consider whether divergences from expectations were due to the implementation or the model itself. We now detail some key problems that were detected and some possible explanations and solutions, the ideal conditions for improving individuality and several interesting results that emerged during testing.

Number of Soft Goals In order to permit behaviour that shows more than one dimensional (simplistic personality) a large number of soft goals (probably at least six) were deemed necessary. However, with n soft goals, for example, the number of contexts that the character can be in is 4^n (given that each soft goal achievement level can take on one of four values: 'H', 'M', 'L', or '-'). This creates a very large number of contexts to be learning about, especially when the reward is non-deterministic. In order to remedy this situation, it seems that the current context should be something other than the soft goals. Some trials were run in which the characters used their current happiness as their context. This resulted in fast learning, but behaviour was too similar and the learning was not situated enough to fulfill the personality requirements of having context-aware behaviour.

Context Ideally, it was thought that the characters would use their soft goals as context so that they would be able to learn that, when the value for soft goal x was low, they would need to do some particular activity to increase the value of that goal. In our implementation, we found that the characters often switched between approximately two contexts during the same time period. This meant that characters were not experiencing the full range of contexts available to them. This could be due to the soft goal achievement level buckets not being complex enough to capture all information. For example, it may be virtually impossible for any character to obtain any more than a "moderate" number of friends, and hence all characters will appear to achieve "make friends" to a level of "M". The cut-offs for the achievement level buckets are based on the entire range that the achievement level can take and the ideal value of the soft goal. So future work, may consider changing the ideal value within the soft goal personality. Another hindrance to character learning in any context in our implementation is the fact that the reward values were found to fluctuate greatly, based on fluctuations in the environment.

Relationship Between Reward and Individuality The results seem to indicate that high individuality can lead to lower personal reward values. For example, Case 4 with 'context off' and Case 2 in both modes show very high individuality and yet very low reward. High individuality, such as in Case 4, means the characters are choosing very different activities at the top-level. However, the different activity choice could be because the characters have not learned how to consistently improve their reward values and are still attempting non-optimal plans. If there is one clear "optimal" path for all characters, then they should all discover this and get higher personal reward values. However, if they have all discovered the same optimal path, then the individuality will decrease. To address these issues, it would be necessary to create more paths that are optimal in the design process. The domain we used was relatively simplistic and therefore the number of possible paths was not as deep as would be likely in a commercial game.

Although it is difficult to draw a conclusion from only five Cases, it seems that if the characters are achieving high reward values, they become more similar and individuality decreases. To confirm this, more Cases with different soft goal personality types within the same Case and across all Cases would need to be used.

Summary Despite these issues, the model does show some promise. Character behaviour does change over time (see Section 5.1.1, page 146); and characters can learn about functional soft goals (see Section 5.1.2, page 152). Also the characters are able to exhibit different behaviour in different contexts (see Section 5.1.3, page 154). Further, we found that characters did choose activities in their own unique manner, based on the individuality measure used. The model was certainly better than random choice of the activities, both in terms of reward and individuality. Although there was no clear distinction between whether the use of contexts improved personal reward and individuality, this may be a consequence of the problems listed above. With further work, it is envisaged that context would emerge as being suitable for many domains, especially domains more complex than the implementation used for our testing.

5.5 Extra Results: Domain-dependent Measures

The results presented so far relate directly to the testing-based research sub-questions and criteria posed in the introductory chapter and are based on the generic model presented in Chapter 3 (page 69). When implementing the characters in our game domain, the characters needed some domain-dependent knowledge, see Section 4.1.2.2 (page 110). This knowledge was divided into opinions and facts. Facts are aspects of this particular world that the character can perceive and do not have a judgement value attached to them, for example, my location[1]. Opinions are facts with a judgement value attached to them, for example, a character may store: 'I like Bec'. In our domain, the characters stored two types of opinions: *attraction towards others*, and personal *happiness*.

Based on *attraction towards others*, characters were able to build friendships that were relatively stable. Although how friendships are developed is based on domain-dependent equations, whether they wanted to make friends was part of their soft goals and therefore somewhat dependent on the stability of the model. The characters also had a *happiness* value that represented how close they were to achieving all of their soft goals.

[1] In some domains, location may need a judgement value attached to it. For example, the character may need to store: 'the location I am in currently is bad'

In this final section of this chapter, we discuss some interesting findings that relate to the opinion-based domain-dependent knowledge. This thesis did not aim to consider the effects of friendship and happiness on our characters. However, these results represent interesting ways that the characters could interact in the social game that was implemented. These results demonstrate the complexity of our model in generating complex characters and indicate possible future avenues of study. The results themselves are presented in Appendix C (page 205) and show graphs of how happiness and friendships fluctuate over time, and friendship networks generated. We now discuss the implications of these extra results.

5.5.1 Happiness

Happiness was found to fluctuate rapidly from very positive to very negative. However, perhaps this is understandable. The characters are rarely achieving their soft goals and if they do suddenly achieve them, it is often for a reason that they seem incapable of learning. Characters are constantly being insulted and have no direct way to stop this behaviour. In fact, it emerged that the potentially best way to avoid being insulted is to insult other people so that the character can make friends, in the hope that these friends will not turn around and insult them. The domain fluctuates because all the characters are all trying to learn what they should be doing. This means that no character can learn a strategy against a particular character because that character may change their strategy as well. However, these problems are problems that we, as humans, face every day. They are also problems that occur in gaming. Human players will not always play consistently and the character needs to be able to adapt to these changes. So, if we can design characters that can cope with constant changes, then they will be able to function better within games. Fluctuating happiness values indicate that the character was changing to match its environment.

Although no players were used for testing, the model (and implementation) are intended for final use in a game for people to play. When watching the game, players are able to observe whether the characters were “happy”, “sad”, or “neutral” based on their smile (for example see the character images in Figure 4.1 on page 105). In our domain, “be happy” could be used as a soft goal based on happiness levels and would then contribute to a character’s context. This soft goal was not used for any of the Cases presented in this thesis. However, if it was used in a future study, the

player would be able to see a factor that relates directly to the character's context. If players cannot see which context (or at least partial context) a character is in when it makes a decision, then the decisions that the characters make may seem erratic or unrealistic. Using happiness may be a potential way to make character behaviour more understandable to a player.

5.5.2 Friendships

We found that characters formed stable friendships, even in their changing domain. This is promising given the difficultly of making friends in our model. Characters could not directly make friends. For example, let us examine how Bec and Gina could become friends. Bec could tell Gina an insult about Heidi, and Gina may like Bec more as a result, but only if Gina does not already like Heidi. Bec was unable to store information on who Gina liked or did not like. If Gina agrees with whatever Bec says, then Bec may like Gina more. However, neither Bec nor Gina can do something 'nice' in order to become friends more quickly. Our implemented game was designed specifically to be difficult to make friends, in order to test the characters and force them to try more ways to make friends and to force them to choose who to be friends with. The (intentional) construction of the domain made it nearly impossible for one character to be friends with everyone; they had to choose. In some Cases (such as Case 1), the characters had no soft goals relating to making friends, so it is not surprising that they only made enemies.

It is very promising that, despite these impediments, the characters were still able to form stable and complex friendships. Further, it is interesting that the complex and relatively stable networks that developed can be related to real world friendships. For example, some of the girls were very popular and formed a sort of 'clique', other girls were happy to be on their own and still other girls were trying to get into the clique (that is, they liked the members of the girls in the clique, but they were not liked by the clique). In our domain, there were only eight characters and yet there still existed a large number of possible friendship networks that could be generated. These networks are dynamic and can change if a new character or a human player enters the environment. It would probably be interesting to investigate how the friendship networks change over longer time periods. In summary, this domain-dependent game

application shows that complex friendships can be easily generated and allow each character's experience to differ from other characters.

5.6 Summary

The results from using the criteria to answer the testing-based research sub-questions can be seen in Table 5.6. Adaptation was found to change behaviour over time, allow characters to learn specific soft goals, and, compared to decisions made using random choice (null hypothesis), adaptation improves both reward and individuality. Characters chose different actions in different contexts. Context was not found to improve individuality and reward in all Cases. Characters were found to be different from each other and, in half the Cases, at least one character was different from the majority of the other characters. Happiness and reward were found to fluctuate greatly over time in response to the dynamic environment. Complex friendship networks were developed showing that our simplistic model was able to generate complex effects. By meeting some of the criteria, the model can be seen to addressing our research goals and questions. We consider the greater meaning of these results in the next chapter as we conclude this thesis.

Research Questions, Testing-based Sub-questions and Criteria for Success	**Criteria Satisfied?**
1. How does adaptation affect character behaviour?	
1.(a) Does behaviour change over time? *Behaviour* changes over time.	Yes
1.(b) Can characters learn about specific, functional goals? When given a functional goal to learn, the majority of characters choose the "correct" action the majority of the time, based on *behaviour*.	Yes
1.(c) How does reward change with time? *Reward* values are on average higher using our model than when random choice is used.	Yes
1.(d) What happens if adaptation is turned off? Compared to when adaptation is turned off; both *individuality* and *reward* are higher when adaptation is used.	Individuality: Yes Reward: Yes
2. How does context affect character behaviour?	
2.(a) Does character behaviour differ in different contexts? For one character's *behaviour*, show that in different contexts the action chosen the majority of times is different.	Yes
2.(b) What happens if context is turned off? Compared to when context is turned off; both *individuality* and *reward* are higher when context is used.	Individuality: Yes in 1/5 Cases Reward: Yes in 1/5 Cases
3. How can personality be implemented so that the same template can be used to create a number of distinct, **individual** characters, according to their behaviour?	
3.(a) Are the behaviours of characters different from each other over time? Character *behaviour* passes the chi-squared test.	Yes
3.(b) Are any individuals obtained? Based on their *individuality*, at least one character is different from the majority of the other characters when they are all based on the same template.	Yes in 2/4 Cases

Table 5.6: Summary of results of testing-based research sub-questions based on criteria for success (as introduced in Table 1.1, page 16). Words in *italics* are the measures of effectiveness.