Preferential Semantics for Causal Fixpoints

Pavlos Peppas[1], Maurice Pagnucco[2], Mikhail Prokopenko[3], and Norman Foo[2]

[1] Knowledge Systems Group, Department of Computing, School of MPCE, Macquarie University, NSW 2109, Australia. Email: pavlos@mpce.mq.edu.au.
[2] Knowledge Systems Group, School of Computer Science and Engineering, University of New South Wales, NSW 2052, Australia. Email: {morri, norman}@cse.unsw.edu.au.
[3] Knowledge Systems Group, CSIRO Mathematical and Information Sciences, Locked Bag 17, North Ryde, NSW 2113, Australia. Email: mikhail.prokopenko@cmis.csiro.au.

Abstract. In this paper we concentrate on the causal theory of action developed by McCain and Turner [2] for computing *ramifications*. Our aim here is to characterise this theory of action in terms of a *preferential-style semantics* in the spirit of Shoham [4]. Such a result would not only place McCain and Turner's theory in a uniform setting, facilitating comparison with other logics of action, but also give a clearer insight into the nature and behaviour of causality captured by their framework. We first show that this objective is not attainable via a traditional preferential semantics. However, preferential semantics is not abandoned entirely. Rather, it is augmented to arrive at the desired result. We maintain that two components — minimal change and causality — are essential in providing a (concise) solution to the frame and ramification problems.

1 Introduction

An early, landmark, development in the field of reasoning about action was Shoham's *preferential semantics* [4]. It provides a semantics for a class of (non-monotonic) logics in which an ordering is imposed over the class of interpretations and the models relating to a particular entailment are subsequently identified as the minimal models under this ordering. Intuitively, the ordering indicates a preference or likelihood over interpretations with only the most preferred being considered in the final outcome (thus sanctioning more inferences than the corresponding classical semantics). The *principle of minimal change* — taking the minimal ("most preferred") models — is the overriding tenet of this approach. Needless to say, work on reasoning about action has progressed but preferential semantics remains an important and intuitive idea.

Currently, a topic of much interest is the role of *causality* in reasoning about action, especially in connection to the *ramification problem*. Over the last few years it has become apparent that *domain constraints* alone are not sufficient to determine the indirect effects of actions and the notion of causality has been suggested as one way of addressing this problem.

In this paper we concentrate on one particular approach that is based on causality, advanced by McCain and Turner [2]. They replace traditional domain constraints by causal associations between formulae $\phi \Rightarrow \psi$ known as *causal laws.* Intuitively, we can think of such a formula as expressing that ϕ *causes* ψ. As such, causal laws express "a relation of determination between the states of affairs that make ϕ and ψ true" [2, p. 1979]. One important aspect of causal laws is that the contrapositive (i.e., $\neg\psi \Rightarrow \neg\phi$) does not hold in general.

The main aim of this paper is to furnish a preferential-style semantics for McCain and Turner's causal theory of action. In particular, the contributions offered are, at least, twofold. Firstly, we show that preferential semantics *per se* cannot capture McCain and Turner's causal theory (Section 4). Secondly, as a result of this, we augment preferential semantics to obtain the desired characterisation. This is achieved by introducing two state selection mechanisms: *state transition systems* (Section 7) which augment the preference structure with a binary relation on states to give the requisite preferential-style semantics and *state elimination systems* (Section 6) which function as an intermediary between causal systems and state transition systems. State elimination systems should not only be viewed as a stepping stone but as providing a clearer picture of the underlying nature of causality.

There are a number of advantages to be gained from these results. They will place causal theories, in the spirit of McCain and Turner's approach, in the context of action theories and (nonmonotonic) logics having a preferential semantics, facilitating a clear comparison between the two. They will give a clearer insight into the nature of causality as captured by these systems. Previous work has shown that minimal change and causality are interreducible under certain conditions (e.g., [1]). One may be tempted to suggest that such a result can be generalised. We maintain that this is not the case but that the two coexist and, in fact, complement each other.

2 Technical Preliminaries

Throughout this paper we will be working with a fixed finitary propositional language $\mathcal{L}$ whose propositional letters we will call *fluents.* The set of all fluents is denoted by $\mathcal{F}$. A *literal* is a fluent or the negation of a fluent. A *state* (or *world*) is defined as a maximal consistent set of literals. The set of all states will be denoted $\mathcal{W}$ and the set of all literals by $\mathcal{N}$. By $[\phi]$ we denote all states consistent with the sentence $\phi \in \mathcal{L}$ (i.e., $[\phi] = \{w \in \mathcal{W} : w \vdash \phi\}$). Occasionally we will refer to $[\phi]$ as ϕ*-states.*

3 Causal Systems

In this section we briefly review the basic concepts underlying McCain and Turner's [2] causal theory of action. This will help to further fix some notation and spell out more clearly the precise objective of this paper.

McCain and Turner introduce a new connective $\Rightarrow$ between sentences φ, ψ of $\mathcal{L}$, and call expressions of the form $\varphi \Rightarrow \psi$ *causal laws* (or *causal rules*).[1] A set of causal laws $\mathcal{D}$ is called a *causal system*. For a set of sentences $\Gamma \subseteq \mathcal{L}$ and a causal system $\mathcal{D}$, the *closure of* Γ *in* $\mathcal{D}$, denoted $C_{\mathcal{D}}(\Gamma)$, is defined as the smallest superset of Γ that is closed under logical implication, and it is such that for any $(\varphi \Rightarrow \psi) \in \mathcal{D}$, if $\varphi \in C_{\mathcal{D}}(\Gamma)$ then $\psi \in C_{\mathcal{D}}(\Gamma)$. We shall say that Γ *causally entails* φ with respect to $\mathcal{D}$, denoted $\Gamma \mid\!\sim_{\mathcal{D}} \varphi$, iff $\varphi \in C_{\mathcal{D}}(\Gamma)$.

One further notion we require is that of a *legitimate state* (wrt a causal system $\mathcal{D}$). A state r is called legitimate with respect to $\mathcal{D}$ if and only if $r = C_{\mathcal{D}}(r) \cap \mathcal{N}$. Intuitively, a state is legitimate if and only if it does not violate any of the causal laws in $\mathcal{D}$. We denote the set of legitimate states (wrt $\mathcal{D}$) by $Legit_{\mathcal{D}}$.

Now McCain and Turner's aim is to determine, given an initial state w and sentence E representing the direct effects (or post-conditions) of an action[2], the set of possible next states $\mathtt{Res}_{\mathcal{D}}(E, w)$. More formally, to any causal system $\mathcal{D}$ we associate a function $\mathtt{Res}_{\mathcal{D}}$ mapping a legitimate state w (the initial state) and sentence E to the set of states $\mathtt{Res}_{\mathcal{D}}(E, w)$ defined as follows:

$$\mathtt{Res}_{\mathcal{D}}(E, w) = \{r \in [E] : r = \{\mathtt{f} : \ (w \cap r) \cup \{E\} \mid\!\sim_{\mathcal{D}} \mathtt{f}\} \ \}$$

Often we shall refer to the members of $\mathtt{Res}_{\mathcal{D}}(E, w)$ as *causal fixpoints*. Intuitively, resultant states are those E-states at which the only changes are the ones that are causally justified.

The aim of this paper can now be seen clearly. Put simply, it is to provide a preferential-style semantics capturing $\mathtt{Res}(E, w)$ for any legitimate state w and sentence E. In other words, to mimic McCain and Turner's fixpoint definition with a process involving a preference ordering imposed over states.

4 Counterexample

At first glance it may seem that all that has to be done to achieve our aims is to determine the restrictions that need to be placed on an ordering $<$ over states to effect the desired result. However, in this section we present an example showing that this straightforward approach will not work in general.

Suppose we are given an initial state $w \in \mathcal{W}$ and a (strict) preference ordering $<_w \subseteq \mathcal{W} \times \mathcal{W}$ on states with the only proviso that $<_w$ satisfy transitivity. In the spirit of preferential semantics [4] we would like to define the states resulting from the occurrence at w of an action with postcondition E, as the minimal E-states wrt $<_w$:

(P) $\quad \mathtt{Res}_{\mathcal{D}}(E, w) = min([E], \ <_w)$

1 For simplicity, throughout this paper we assume that the antecedent of any causal law is consistent.

2 Since actions play no immediate role in this framework we shall not refer to them explicitly but only through their direct effects.

In McCain and Turner's approach it is not necessary to consider illegitimate states as possible resultant states since they violate the causal laws. Therefore, we first concentrate on the following variant of (P).

$$(\mathrm{P}') \qquad \mathtt{Res}_{\mathcal{D}}(E, w) = min([E],\ <_w) \cap\ Legit_{\mathcal{D}}$$

The first result shows that (P′) cannot be satisfied in general.

Theorem 1. *There exists a causal system $\mathcal{D}$ and (initial) state w such that no ordering on states satisfies* (P′).

Proof: Assume that $\mathcal{L}$ has five propositional letters $\mathtt{a},\ \mathtt{b},\ \mathtt{c},\ \mathtt{d},\ \mathtt{e}$. Let the initial state be $w = \{\mathtt{a},\ \mathtt{b},\ \mathtt{c},\ \mathtt{d},\ \mathtt{e}\}$ and define s_1, s_2 and s_3 to be the following states: $s_1 = \{\neg\mathtt{a},\ \neg\mathtt{b},\ \mathtt{c},\ \mathtt{d},\ \neg\mathtt{e}\}$, $s_2 = \{\mathtt{a},\ \neg\mathtt{b},\ \neg\mathtt{c},\ \mathtt{d},\ \neg\mathtt{e}\}$ and $s_3 = \{\mathtt{a},\ \neg\mathtt{b},\ \mathtt{c},\ \neg\mathtt{d},\ \mathtt{e}\}$. Finally let $\mathcal{D}$ be the following causal system: $\mathcal{D} = \{\bigwedge s_i \Rightarrow\ \ false$ (for $s_i \neq w,\ s_1,\ s_2,\ s_3$), $\neg\mathtt{a} \Rightarrow \neg\mathtt{b},\ \ \neg\mathtt{d} \Rightarrow \neg\mathtt{b}\}$.

Consider now the following direct effects (post-conditions) of actions. $\Delta_1 = (\neg\mathtt{a} \vee \neg\mathtt{c}) \wedge \neg\mathtt{e}$, $\Delta_2 = (\neg\mathtt{a} \vee \neg\mathtt{d})$ and $\Delta_3 = (\bigwedge s_2) \vee (\bigwedge s_3)$. Clearly, state s_1 satisfies Δ_1 and Δ_2; s_2 satisfies Δ_1 and Δ_3; and, s_3 satisfies Δ_2 and Δ_3.

Suppose a (transitive) ordering on states $<_w$ satisfying condition (P′) exists. Now, the following is easily (albeit tediously) verified. $\mathtt{Res}_{\mathcal{D}}(\Delta_1,\ w) = s_1$ from which we conclude $s_1 <_w s_2$. $\mathtt{Res}_{\mathcal{D}}(\Delta_2,\ w) = s_3$, therefore $s_3 <_w s_1$. By the transitivity of $<_w$ we infer that $s_3 <_w s_2$. However, $\mathtt{Res}_{\mathcal{D}}(\Delta_3,\ w) = \{s_2,\ s_3\}$ from which it follows that $s_3 \not<_w s_2$ contradicting the above. ■

Returning our attention to (P) which is a closer rendering of preferential semantics than (P′), the result above can now be used to show Theorem 2 below.[3] We therefore conclude that a traditional preferential semantics (characterised by condition (P)) cannot, in general, be supplied for McCain and Turner's causal theory of action.

Theorem 2. *There exists a causal system $\mathcal{D}$ and (initial) state w such that no ordering on states satisfies* (P).

This will not lead us to abandon preferential semantics entirely. Our strategy will be to retain as much of the spirit of preferential semantics as possible and augment it with a further mechanism that takes into account the presence of causality.

5 Three State-Selection Mechanisms

From a purely technical point of view, McCain and Turner's causal theory of action can be regarded as nothing more than a state-selection mechanism. More precisely, given some background knowledge expressed in the form of *causal laws*, McCain and Turner provide a method for selecting, for any state w and any sentence E representing the initial state and direct effects respectively, a

[3] Proofs will generally be omitted due to lack of space.

subset $\mathtt{Res}_{\mathcal{D}}(E, w)$ of $[E]$. The intention is of course that $\mathtt{Res}_{\mathcal{D}}(E, w)$ contains precisely those states that can possibly result from the occurrence of an action with post-condition E at the initial state w.

However, we can adopt a more abstract perspective and view it simply as a selection function. On this view $\mathtt{Res}_{\mathcal{D}}(E, w)$ is just a function selecting the "best" states from among E (wrt w). In this paper, we now proceed as follows. Firstly, in addition to McCain and Turner's selection mechanism, we present two other types of systems that provide means for selecting states which we call respectively *state elimination systems* and a *state transition systems*. State transition systems constitute the preferential-like semantics characterising McCain and Turner's causal theory of action.

The desired characterisation is achieved in two steps: we first show how to translate from causal rules to state elimination systems (while preserving the selection function). Then we show how (under certain conditions) we can move from state elimination systems to state transition systems (again in a way that preserves the selection function). Let us now take a more detailed look at each of these systems.

6 State Elimination Systems

The performing of an action with known direct effects E makes certain states — the E-states (or a subset of them) — entertainable as possible resultant states. The basic idea behind state elimination systems is the use of certain elimination rules to reject states from further consideration. Any state rejected through the application of an elimination rule violates a causal relationship deemed to hold in the world.

Definition 3. (State elimination rule)
A *state elimination rule* (or simply, *elimination rule*) is an expression of the form $\{r_1, r_2, \cdots, r_k, r_{k+1}, \cdots, r_n\} \rhd \{r_1, r_2, \cdots, r_k\}$ where each r_i is a state.

Intuitively a state elimination rule works as follows. Assume that, according to your current beliefs, the states you consider possible are among $\{r_1, \cdots, r_n\}$. Then the above rule allows you to eliminate or disregard the worlds $r_{k+1}, \cdots, r_n$. A state elimination system $\mathcal{S}$ is a set of elimination rules.

In a state elimination system we are working with a set of states; the states currently being entertained. Applying elimination rules to such a set of states potentially allows us to rule out illegitimate states and concentrate our attention on the remaining ones. All elimination rules would, of course, need to be applied in turn in order to ensure that no illegitimate states are considered possible.

Definition 4. ($\leadsto$ and $\overset{*}{\leadsto}$)
Given a state elimination system $\mathcal{S}$, we shall say that a set of states Q *yields in one step* a set of states R, which we denote by $Q \leadsto R$, iff there is an elimination rule $X \rhd Y$ such that $Q \subseteq X$ and $R = Q \cap Y$. We define $\overset{*}{\leadsto}$ to be the reflexive transitive closure of $\leadsto$.

At some point the set of entertained states will reach a point of *equilibrium*. That is, the application of any further rules will no longer eliminate states.

Definition 5. (Compact state)
We shall say that a set of states Q is *compact* (in S) iff for any R such that $Q \overset{*}{\leadsto} R$, $Q = R$. If Q is a singleton and compact, we will also call the state in Q compact.

Before presenting the state selection mechanism associated with elimination systems, we need one more definition.

Definition 6. (E-predecessor)
For any two states w, r and any sentence E, we define the E-predecessors of r with respect to w, to be the set $\langle\!| r, E |\!\rangle_w = \{r' : r' \in [E] \text{ and } \mathit{Diff}(w, r') \subseteq \mathit{Diff}(w, r)\}$ (where $\mathit{Diff}(w, r)$ represents the symmetric difference of r and w (i.e., $(r \setminus w) \cup (w \setminus r)$) as in the PMA [5]).

Clearly any $r \in [E]$ is an E-predecessor of itself with respect to w (i.e., $r \in \langle\!| r, E |\!\rangle_w$). The E-predecessors of r with respect to w are simply those E-states which agree with w on at least those fluents where w and r agree and possibly others. If one considers a PMA ordering [5] of states $\leq_W$ (i.e., $r \leq_W s$ iff $\mathit{Diff}(w, r) \subseteq \mathit{Diff}(w, s)$), then the E-predecessors are simply those E-states closer to w than r.

We now define the state selection mechanism at the heart of state elimination systems.

Definition 7. ($\mathtt{Next}_{\mathcal{S}}(E, w)$)
To any state elimination system $\mathcal{S}$ we associate a function $\mathtt{Next}_{\mathcal{S}}$ (mapping a compact (in $\mathcal{S}$) state w and a sentence E to the set of states $\mathtt{Next}_{\mathcal{S}}(E, w)$) defined as follows: $\mathtt{Next}_{\mathcal{S}}(E, w) = \{r \in [E] : r \text{ is compact (in } \mathcal{S}) \text{ and } \langle\!| r, E |\!\rangle_w \overset{*}{\leadsto} \{r\}\}$

As we shall see in section 6.1, $\mathtt{Next}_{\mathcal{S}}(E, w)$ will be used to capture $\mathtt{Res}_{\mathcal{D}}(E, w)$ in terms of state elimination systems. Intuitively, r is a permissible next state if all its E-predecessors with respect to w are eliminated by the rules in $\mathcal{S}$ but r <u>alone</u> survives. If r is eliminated then, as indicated above, it must violate some causal law. On the other hand, if r survives together with some other state(s), then it means that there is some ambiguity for which causality cannot account. More to the point, since our intention is not to abandon the principle of minimal change entirely but merely to augment it with causality, then it means that there is some closer state (i.e., a state admitting "less change") which is consistent with the causal laws being considered. This is precisely why we need only consider the E-predecessors of a state r when determining its membership in $\mathtt{Next}_{\mathcal{S}}(E, w)$. We want to know whether a state is illegitimate or whether there is a "closer" state satisfying the causal forces at work. If so, we need no longer entertain it as a permissible resultant state.

6.1 Causal Systems and State Elimination Systems

In this section we define mappings from causal systems to state elimination systems and vice versa. This will serve to identify the interrelationship between the two. For this purpose we require the following definition.

Definition 8. (Selection-equivalent)
We shall say that a causal system $\mathcal{D}$ is *selection-equivalent* to a state elimination system $\mathcal{S}$ iff $\texttt{Res}_{\mathcal{D}}(E,w) = \texttt{Next}_{\mathcal{S}}(E,w)$, for every sentence E and state w.

The term selection-equivalent will also be used to relate (in the obvious way) not just a causal system and a state elimination system, but any two state selection mechanisms (for example, a causal system and a state transition system, two different state elimination systems, etc.).

The main result of this section — an exact characterisation of McCain and Turner's causal theory of action in terms of state elimination systems – can now be established.

Theorem 9. *For every causal system there is a selection-equivalent state elimination system. Conversely, for every state elimination system there is a selection-equivalent causal system.*

Proof (Sketch)
($\Rightarrow$)
Let $\mathcal{D}$ be an arbitrary causal system. For every causal law $\varphi \Rightarrow \psi$ in $\mathcal{D}$, produce the elimination rule $[\varphi] \rhd [\varphi \wedge \psi]$. Call $\mathcal{S}$ the set of elimination rules so produced. It is not hard to verify that for any legitimate state w and sentence E, $\texttt{Res}_{\mathcal{D}}(E,w) = \texttt{Next}_{\mathcal{S}}(E,w)$ (simply notice that for any state r, $[(w \cap r) \cup \{E\}] = \langle\!\langle r, E \rangle\!\rangle_w$).
($\Leftarrow$)
Let $\mathcal{S}$ be an arbitrary state elimination system. For every elimination rule $X \rhd Y$ produce the causal law $\varphi \Rightarrow \psi$, where φ, ψ are such that $[\varphi] = X$ and $[\psi] = Y$ (since our language is a finitary propositional one, such φ and ψ always exist). The set of causal laws so produced, call it $\mathcal{D}$, is selection-equivalent to $\mathcal{S}$. ■

Note, in this proof, the relationship between causal laws and elimination rules: $\phi \Rightarrow \psi$ if and only if $[\phi] \rhd [\phi \wedge \psi]$ (or, equivalently, $[\phi] \rhd [\phi] \cap [\psi]$).

The following two definitions will be important for the next section but we introduce them here in order to highlight a relationship that exists between a certain class of state elimination systems and a class of causal systems.

Definition 10. ($\mathcal{S}$ Closed under union)
We shall say that a state elimination system $\mathcal{S}$ is *closed under union* iff whenever $(X \rhd Y)$ and $(Q \rhd R)$ are in $\mathcal{S}$ then $(X \cup Q) \rhd (Y \cup R)$ is also in $\mathcal{S}$.

Definition 11. ($\mathcal{D}$ Closed under disjunction)
We shall say that a causal system $\mathcal{D}$ is *closed under disjunction* iff whenever $\varphi \Rightarrow \psi$ and $\chi \Rightarrow \zeta$ are in $\mathcal{D}$ then $(\varphi \vee \chi) \Rightarrow (\psi \vee \zeta)$ is also in $\mathcal{D}$.

The relationship is given by the following result.

Theorem 12. *Every causal system closed under disjunction is selection-equivalent to a state elimination system closed under union. Conversely, every state elimination system closed under union is selection-equivalent to a causal system closed under disjunction.*

To see this rather quickly, note that disjunction in terms of sentences ($\phi \vee \psi$) corresponds to union in terms of states ($[\phi \vee \psi] = [\phi] \cup [\psi]$).

7 State Transition Systems

A state transition system consists of a binary relation over states which, as the name suggests, represents transitions between states. We group together some definitions required in the context of such systems.

Definition 13. A *state transition system* $\mathcal{M}$ is a binary relation on the set $\mathcal{W}$ of states (i.e., $\mathcal{M} \subseteq \mathcal{W} \times \mathcal{W}$). Whenever $\langle r, r' \rangle \in \mathcal{M}$ we will write $r \rightharpoonup r'$. We define $\overset{*}{\rightharpoonup}$ to be the reflexive transitive closure of $\rightharpoonup$. Whenever $r \overset{*}{\rightharpoonup} r'$ we shall say that r' is *reachable* from r. We shall say that a state r is *final* (in $\mathcal{M}$) iff for any r' such that $r \overset{*}{\rightharpoonup} r'$, $r = r'$.

Intuitively, $\mathcal{M}$ represents transitions between states due to the presence of causality.

We shall now define the notion of *successor states* $\mathtt{Succ}_{\mathcal{M}}(E, w)$ for a state transition system $\mathcal{M}$, initial state w and direct effects E of an action. This identifies the resultant states and it is this that will be used to capture $\mathtt{Res}_{\mathcal{D}}(E, w)$ (and $\mathtt{Next}_{\mathcal{S}}(E, w)$).

Definition 14. ($\mathtt{Succ}_{\mathcal{M}}(E, w)$)
To any state transition system $\mathcal{M}$ we associate a function $\mathtt{Succ}_{\mathcal{M}}$ (mapping a final (in $\mathcal{M}$) state w and a sentence E to the set of states $\mathtt{Succ}_{\mathcal{M}}(E, w)$) defined as follows: $\mathtt{Succ}_{\mathcal{M}}(E, w) = \{r' \in [E] : r'$ is final (in $\mathcal{M}$) and $r \overset{*}{\rightharpoonup} r'$ for all $r \in \langle\!\!\langle r', E \rangle\!\!\rangle_w\}$.

There are some important points to note about $\mathtt{Succ}_{\mathcal{M}}(E, w)$ that are at the very heart of our preferential-style semantics. Notice that $\mathtt{Succ}_{\mathcal{M}}(E, w)$ relies on two very important components. Firstly, a preference ordering over states used to determine the E-predecessors of r' (wrt w — i.e., $\langle\!\!\langle r', E \rangle\!\!\rangle_w$). In fact, it is a PMA ordering as previously noted. Secondly, the binary relation $\mathcal{M}$ on states. *The former embodies the principle of minimal change while the latter captures the notion of causality.* Again, only the E-predecessors of r' (wrt w) need be considered due to our desire to adhere to the principle of minimal change as much as possible. Also note that a resultant state must be reachable via "causality" from all of its E-predecessors. If it were not, then there must be a state "closer" to w consistent with the causal laws and the minimality/causality duo would be violated.

7.1 State Elimination Systems and State Transition Systems

In this section we obtain mappings between state elimination systems and state transition systems and vice versa. Then, using the results of section 6.1, we obtain a correspondence between causal systems and state transition systems; thus a preferential-style semantics for McCain and Turner's causal theory of actions.

Theorem 15. *For every state elimination system $\mathcal{S}$ that is closed under union, there exists a selection-equivalent state transition system $\mathcal{M}$. Conversely, for every state transition system $\mathcal{M}$ there is a selection-equivalent state elimination system $\mathcal{S}$ that is closed under union.*

Proof. (Sketch)

$(\Rightarrow)$

Let $\mathcal{S}$ be a state elimination system closed under union. Let $\mathcal{S}'$ be a unary state elimination system that is selection-equivalent to $\mathcal{S}$. From $\mathcal{S}'$ we shall construct a selection-equivalent state transition system $\mathcal{M}$. First however we need some definitions.

Consider an arbitrary set of states Q with cardinality $n+1$. We shall say that the string of elimination rules $\sigma_1; \sigma_2; \cdots; \sigma_n$ *dissolves* Q iff after applying these rule successively (in that order), all but one of the states of Q are eliminated and, moreover, the one remaining state is compact.

Assume that $\sigma_1; \sigma_2; \cdots; \sigma_n$ dissolves Q and for all $1 \leq i \leq n$, let r_i be the state of Q that is eliminated by the rule σ_i; let us also call w the one state of Q that is not eliminated. We shall call the sequence $r_1; r_2; \cdots; r_n; w$ a *trace* for Q (in $\mathcal{S}'$).

We now construct from $\mathcal{S}'$ a state transition system $\mathcal{M}$ as follows. For any two states r, r', $\langle r, r' \rangle \in \mathcal{M}$ iff there is a dissolvable set of states Q containing r and r', such that for some trace of Q in $\mathcal{S}'$, r' appears immediately after r. It can be shown that $\mathcal{M}$ is selection-equivalent to $\mathcal{S}'$.

$(\Leftarrow)$

Proved by essentially reversing the construction presented above. ■

By combining theorems 12 and 15 we derive the following corollary, which is the central result of this article.

Corollary 16. *For every causal system $\mathcal{D}$ closed under disjunction, there exists a selection-equivalent state transition system $\mathcal{M}$. Conversely, for every state transition system $\mathcal{M}$ there exists a selection-equivalent causal system $\mathcal{D}$ which is closed under disjunction.*

This says that our preferential-style semantics in terms of state transition systems — consisting of a preference ordering (in this case, PMA ordering) over states augmented by a binary relation on states — exactly characterises McCain and Turner causal systems where the causal laws are closed under disjunction.

8 Conclusion and Future Work

In this article we provided an augmented preferential semantics for McCain and Turner's causal theory of action. In doing so, we first demonstrated that traditional preferential semantics using the principle of minimal change cannot achieve this task alone. To remedy this we did not abandon minimal change but rather we enhanced preferential semantics with a binary relation capturing the force of causality. That is, causal laws were transformed (via state elimination systems) into a binary relation on states. Moreover, we maintain that these two components – minimal change and causality – are both essential in providing a (concise) solution to the frame and ramification problems.

Placing McCain and Turner's causal theory of action in the uniform setting of preferential-style semantics facilitates comparison with other logics of action. This is clearly an interesting avenue for future work. Another equally important research direction is a comparison between the state transition semantics introduced herein, and the *causal propagation semantics* proposed by Sandewall [3]. Establishing a link between the two will indirectly connect (by means of the results reported herein) Sandewall's semantics with McCain and Turner's theory of action, thus giving a clearer insight into the nature of causality captured by these various formalisms.

References

1. Joakim Gustafsson and Patrick Doherty. Embracing occlusion in specifying the indirect effects of actions. In L. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on Knowledge Representation and Reasoning*. Morgan-Kaufmann, 1996.
2. Norman McCain and Hudson Turner. A causal theory of ramifications and qualifications. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1978–1984. Montreal, 1995.
3. Erik Sandewall. Assessments of ramification methods that use static domain constraints. In L. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on Knowledge Representation and Reasoning*. Morgan-Kaufmann, 1996.
4. Yoav Shoham. *Reasoning About Change*. MIT Press, Cambridge, Massachusetts, 1988.
5. Marianne Winslett. Reasoning about actions using a possible models approach. In *Proceedings of the Seventh National Artificial Intelligence Conference*, San Mateo, CA., 1988. Morgan Kaufmann Publishers.