

Session 3

In today's session, we will tackle the second half of the representation theorem introduced last time: the “completeness” part. Specifically, we will prove that if a revision operator $*$ satisfies the six basic and two supplementary postulates, then for any belief set B , we can construct a corresponding ordering \preceq_B . This ordering will satisfy the four structural properties outlined [in the previous session](#) and guarantee that:

$$B * \phi = T(\min_B([\phi]))$$

I have pasted the full statement of the theorem below as a refresher. In the following section, we will dive directly into the proof.

Theorem 5 (Representation Theorem for Belief Revision). \checkmark

Let \mathcal{L} be the propositional language defined previously and W be the corresponding space of possible worlds. Let $B \subseteq \mathcal{L}$ be a belief set.

(1) If \preceq_B is a plausibility ordering on W that satisfies Connectedness, Transitivity, Centeredness, and the Limit Assumption, then the revision operator $*$ defined by:

$$B * \phi := T(\min_B([\phi]))$$

satisfies all basic and supplementary AGM postulates.

(2) Assume Φ is finite. If $*$ is a revision operator on B that satisfies all eight AGM postulates, then there exists a plausibility ordering \preceq_B on W satisfying Connectedness, Transitivity, Centeredness, and the Limit Assumption, such that for every formula $\phi \in \mathcal{L}$:

$$B * \phi = T(\min_B([\phi]))$$

Part Two: “Completeness” (Finite Language)

Introduction

Let us now prove the second half of the [Representation Theorem](#)—the “Completeness” direction. This result establishes that any belief revision operator $*$ satisfying the eight AGM postulates discussed in [Session 1](#) can be reconstructed using our possible worlds semantics.

To proceed with this proof, **we must now explicitly assume that our propositional language \mathcal{L} is finite**. This assumption allows us to construct formulas that isolate specific sets of possible worlds.

Let $\Phi = \{p_1, \dots, p_n\}$ be the finite set of propositional variables generating \mathcal{L} . For any world $w \in W$, we define the **state description** δ_w as the conjunction of all literals satisfied by w :

$$\delta_w := \bigwedge_{p \in \Phi} l_p \quad \text{where } l_p = \begin{cases} p & \text{if } w \models p \\ \neg p & \text{if } w \not\models p \end{cases}$$

We define the function $form : \mathcal{P}(W) \rightarrow \mathcal{L}$ such that for any subset of worlds $V \subseteq W$:

$$form(V) := \bigvee_{w \in V} \delta_w$$

For a single world w , we simply write $form(w)$ instead of $form(\{w\})$.

 **A Technical Detail on $form$** \checkmark

Strictly speaking, defining $form(V)$ as simply $\bigvee_{w \in V} \delta_w$ does not yield a mathematically well-defined function. Syntactically, the formula $p_1 \vee \neg p_2$ is a completely different string than $\neg p_2 \vee p_1$, even though they are logically and semantically identical. To make $form$ a proper function that maps to a **unique formula**, we would need to specify a strict lexicographical ordering for how the disjuncts are arranged.

Because the AGM postulates guarantee that our revision operator treats logically equivalent formulas identically (Postulate 6), we can safely abstract away from these syntactic details for the remainder of the proof.

To see how this works, consider a language \mathcal{L} generated by the propositional variables $\Phi = \{a, b, c, d\}$. Recall that a possible world for \mathcal{L} is a valuation function $w : \{a, b, c, d\} \rightarrow \{0, 1\}$. Consider the possible world w such that:

1. $w(a) = 1$
2. $w(b) = 0$
3. $w(c) = 1$
4. $w(d) = 0$

(Recall that $w(a) = 1$ is strictly equivalent to saying $w \models a$, and $w(a) = 0$ is equivalent to saying $w \not\models a$.)

Thus, w satisfies a and c , and falsifies b and d . Recall that δ_w is the conjunction of literals l_p , for all $p \in \Phi$, such that $l_p = p$ iff $w \models p$, and $l_p = \neg p$ iff $w \not\models p$. Therefore, we have that:

$$\delta_w = a \wedge \neg b \wedge c \wedge \neg d$$

This conjunction is the **state description of w** .

Now, consider the set of worlds $\{w, w'\}$, where w' is identical to w except that it satisfies b (i.e., $w'(b) = 1$). The state description for w' is therefore $\delta_{w'} = a \wedge b \wedge c \wedge \neg d$.

By definition, $form(\{w, w'\})$ is equal to the disjunction of their individual state descriptions:

$$(a \wedge \neg b \wedge c \wedge \neg d) \vee (a \wedge b \wedge c \wedge \neg d)$$

Because the variable b takes opposite truth values across the two disjuncts while all other variables remain constant, this is logically equivalent to the much simpler and more readable formula:

$$a \wedge c \wedge \neg d$$

Now that we have clarified the mechanics of δ_w and $form(\cdot)$, before proving the completeness part of the [Representation Theorem](#), let me state a crucial lemma that we will need later on.

Lemma 1.

Let \mathcal{L} be finite and let $form$ be the function defined above.

$$\text{For all } w' \in W : w' \models form(w) \iff w' = w$$

Proof. Let $w \in W$. By definition, $form(w)$ is the conjunction of all literals satisfied by w . Let $w' \in W$ be an arbitrary world.

- (\Leftarrow) If $w' = w$, w' assigns the exact same truth values to all propositional variables as w . Hence, it

satisfies every literal in $form(w)$, meaning $w' \models form(w)$.

- (\implies) If $w' \neq w$, there must be at least one propositional variable $p \in \Phi$ on which they differ. Without loss of generality, suppose $w \models p$ and $w' \not\models p$. Then the literal p is a conjunct in $form(w)$. Since $w' \not\models p$, w' fails to satisfy this conjunct, and therefore $w' \not\models form(w)$.

Thus, $w' \models form(w)$ if and only if $w' = w$. □

In other words, Lemma [Lemma 1](#) shows that the formula $form(w)$ *uniquely characterizes* the world w . There exists exactly one possible world in W that satisfies $form(w)$: w itself. (Differently but equivalently put, any world that satisfies $form(w)$ must be identical to w).

This lemma has a very important corollary, which will play a crucial role in our completeness proof.

Corollary 2.

Let \mathcal{L} be finite, let $form$ be the function defined above, and let $[\cdot]$ be the [truth-set function](#).

$$\text{For all } V \subseteq W : [form(V)] = V$$

Proof. We prove this equality by mutual inclusion. Recall that by definition, $form(V) = \bigvee_{w \in V} form(w)$.

- (\subseteq) Suppose $w' \in [form(V)]$. By the definition of the truth-set function, $w' \models \bigvee_{w \in V} form(w)$. By the semantics of disjunction, there exists at least one world $w \in V$ such that $w' \models form(w)$. By Lemma [Lemma 1](#), this strictly entails $w' = w$. Since $w \in V$, it follows that $w' \in V$.
- (\supseteq) Suppose $w' \in V$. By Lemma [Lemma 1](#), we know that $w' \models form(w')$. Because $form(w')$ is one of the disjuncts in the formula $\bigvee_{w \in V} form(w)$, classical semantics guarantees that $w' \models \bigvee_{w \in V} form(w)$. Thus, $w' \models form(V)$, which means $w' \in [form(V)]$.

Therefore, $[form(V)] = V$. □

In other words, for a finite language, $form(\cdot)$ and $[\cdot]$ (i.e., the truth-set function) act as **inverses of each other up to logical equivalence**. By this, we mean that while translating a set of worlds into a formula and back returns the exact same set of worlds (i.e., $[form(V)] = V$), translating a formula into its truth-set and back yields a formula that is *logically equivalent* to the original, even if it is not necessarily the exact same syntactic string (i.e., $form([\phi]) \dashv\vdash \phi$, though generally $form([\phi]) \neq \phi$). (Find an example to convince yourself that this is true.) These functions thus allow us to translate between subsets of worlds and syntactic formulas without losing any semantic information.

Note that this holds *only* for a finite language. If \mathcal{L} were an infinite language, δ_w would not be a well-formed formula, as its construction would require **infinitely long conjunctions**, which our syntax prohibits.

(N.b., a similar point was made in [Lemma 4 \(point 3\)](#) of the previous session, where it is shown that $T(\cdot)$ and $[\cdot]$ act as inverses of each other if and only if \mathcal{L} is a finite language.)

The Completeness Proof

Lemma 3 (Completeness).

Let \mathcal{L} be a finite propositional language and W be the corresponding space of possible worlds. Let $B \subseteq \mathcal{L}$ be a belief set. If $*$ is a revision operator on B that satisfies all eight AGM postulates, then there exists a plausibility ordering \preceq_B on W satisfying Connectedness, Transitivity, Centeredness, and the Limit Assumption, such that for every formula $\phi \in \mathcal{L}$:

$$B * \phi = T(\min_B([\phi]))$$

Proof. Assume $B \subseteq \mathcal{L}$ is a belief set, and let $*$ be a revision operator on B that satisfies all eight AGM postulates. Because \mathcal{L} is finite, we can define the weak plausibility ordering \preceq_B on W directly from the revision operator using the exact method of Katsuno and Mendelzon (1991). For any $w_1, w_2 \in W$:

$$w_1 \preceq_B w_2 \quad : \iff \quad w_1 \in [B] \text{ or } w_1 \in [B * \text{form}(\{w_1, w_2\})]$$

Now, let us prove that Connectedness, Transitivity, Centeredness, and the Limit Assumption hold for \preceq_B .

Note >

In fact, K&M's definition of \preceq_B is redundant, for it could have been the following:

$$w_1 \preceq_B w_2 \quad : \iff \quad w_1 \in [B * \text{form}(\{w_1, w_2\})]$$

The reason is that, if $w_1 \in [B]$, then by Vacuity $w_1 \in [B * \text{form}(\{w_1, w_2\})]$.

Proof. Suppose that $w_1 \in [B]$, that \preceq_B is defined as above, and that $*$ satisfies AGM's postulates for revision. Then, $B \not\vdash \text{form}(\{w_1, w_2\})$, i.e. there exists a possible world that is both in $[B]$ and $[\text{form}(\{w_1, w_2\})] = \{w_1, w_2\}$ (by [Corollary 2](#)), i.e. w_1 . By the Vacuity, it follows that $B * \text{form}(\{w_1, w_2\}) = \text{Cn}(B \cup \{\text{form}(\{w_1, w_2\})\})$. In set-theoretic terms, it means that

$$\begin{aligned} [B * \text{form}(\{w_1, w_2\})] &= [B] \cap [\text{form}(\{w_1, w_2\})] \\ &= [B] \cap \{w_1, w_2\} \end{aligned}$$

Recall that $w_1 \in [B]$ and obviously $w_1 \in \{w_1, w_2\}$. Therefore, $w_1 \in [B * \text{form}(\{w_1, w_2\})]$. □

Despite this result, I will nonetheless use K&M's original, redundant version.

I. Part One: Properties of \preceq_B .

1. Connectedness. We must show that for any $w_1, w_2 \in W$, either $w_1 \preceq_B w_2$ or $w_2 \preceq_B w_1$.

If either $w_1 \in [B]$ or $w_2 \in [B]$, then the condition is trivially satisfied by definition.

Suppose neither world is in $[B]$. Let $\psi = \text{form}(\{w_1, w_2\})$. Since $[\psi] = \{w_1, w_2\} \neq \emptyset$, Postulate 5 (Consistency) ensures that $[B * \psi] \neq \emptyset$. By Postulate 2 (Success), $B * \psi \vdash \psi$, which semantically means $[B * \psi] \subseteq [\psi] = \{w_1, w_2\}$. Because $[B * \psi]$ is a non-empty subset of $\{w_1, w_2\}$, it must contain w_1 , w_2 , or both.

1. If $w_1 \in [B * \text{form}(\{w_1, w_2\})]$, then $w_1 \preceq_B w_2$.
2. Suppose that $w_2 \in [B * \text{form}(\{w_1, w_2\})]$. Note that $\text{form}(\{w_1, w_2\})$ may be distinct but is logically equivalent to $\text{form}(\{w_2, w_1\})$. By Postulate 6 (Congruence), logically equivalent formulas yield identical revised belief sets. Thus, $w_2 \in [B * \text{form}(\{w_2, w_1\})]$, meaning $w_2 \preceq_B w_1$.

Therefore, \preceq_B is connected.

2. Transitivity. Suppose $w_1 \preceq_B w_2$ and $w_2 \preceq_B w_3$. We must show $w_1 \preceq_B w_3$.

Case 1: $w_1 \in [B]$. By definition, $w_1 \preceq_B w_3$ immediately follows.

Case 2: $w_1 \notin [B]$. We need to prove that $w_1 \in [B * \text{form}(\{w_1, w_3\})]$.

First, since $w_1 \preceq_B w_2$ and $w_1 \notin [B]$, our definition requires $w_1 \in [B * \text{form}(\{w_1, w_2\})]$.

Second, notice that we must also have $w_2 \notin [B]$. Suppose to the contrary that $w_2 \in [B]$. First, it immediately follows that $[B] \cap [form(\{w_1, w_2\})] \neq \emptyset$, for by [Corollary 2](#) $[form(\{w_1, w_2\})] = \{w_1, w_2\}$. This entails that $B \not\vdash \neg form(\{w_1, w_2\})$. By Postulate 4 (Vacuity), it follows that:

$$B * form(\{w_1, w_2\}) = Cn(B \cup \{form(w_1, w_2)\})$$

which implies:

$$[B * form(\{w_1, w_2\})] = [Cn(B \cup \{form(w_1, w_2)\})]$$

Now, by [Session 2 – Lemma 2](#), it follows that:

$$[B * form(\{w_1, w_2\})] = [B] \cap \{w_1, w_2\}$$

Since $w_1 \notin [B]$ and, *ex hypothesi*, $w_2 \in [B]$, it follows that $[B * form(\{w_1, w_2\})] = [B] \cap \{w_1, w_2\} = \{w_2\}$, which contradicts $w_1 \in [B * form(\{w_1, w_2\})]$. Therefore, $w_2 \notin [B]$.

Thirdly, since $w_2 \preceq_B w_3$ and $w_2 \notin [B]$, it must be that $w_2 \in [B * form(\{w_2, w_3\})]$. Again, it must be the case that $w_3 \notin [B]$. (We can show in a similar way as above that, if $w_3 \in [B]$, a contradiction follows.)

Now that we have collected many pieces of information regarding w_1, w_2, w_3 :

1. $w_1 \preceq_B w_2$ and $w_2 \preceq_B w_3$,
2. $w_1, w_2, w_3 \notin [B]$, and
3. $w_1 \in [B * form(\{w_1, w_2\})]$ and $w_2 \in [B * form(\{w_2, w_3\})]$.

So, we can finally prove the statement we need, i.e. $w_1 \in [B * form(\{w_1, w_3\})]$. To do so, we proceed as follows. First, we let $\psi = form(\{w_1, w_2, w_3\})$. Then, we show that, since $w_1 \preceq_B w_2 \preceq_B w_3$, it must be the case that $w_1 \in [B * \psi]$. Then, since obviously $w_1 \in [B * \psi] \cap \{w_1, w_3\}$, where $\{w_1, w_3\} = [form(\{w_1, w_3\})]$, we can apply the supplementary postulates to show that $w_1 \in [B * form(\{w_1, w_3\})]$.

First of all, since $[\psi] \neq \emptyset$, by Postulate 5 (Consistency) $[B * \psi] \neq \emptyset$. By Postulate 2 (Success), $[B * \psi] \subseteq \{w_1, w_2, w_3\}$. In other words, $[B * \psi]$ must be a non-empty subset of $\{w_1, w_2, w_3\}$.

We evaluate which worlds are in $[B * \psi]$ via three exhaustive subcases:

Subcase 2.1: $w_2 \in [B * \psi]$. This means $[B * \psi] \cap [form(\{w_1, w_2\})] \neq \emptyset$. This means that the antecedent of Postulate 8 holds, i.e. that $B * \psi \not\vdash \neg form(\{w_1, w_2\})$. By applying both Postulate 7 (Superexpansion) and Postulate 8 (Subexpansion), we obtain the following:

$$B * (\psi \wedge form(\{w_1, w_2\})) = Cn(B * \psi \cup \{form(\{w_1, w_2\})\})$$

This immediately yields:

$$[B * (\psi \wedge form(\{w_1, w_2\}))] = [Cn(B * \psi \cup \{form(\{w_1, w_2\})\})]$$

and [Session 2 – Lemma 2](#) yields the following:

$$[B * (\psi \wedge form(\{w_1, w_2\}))] = [B * \psi] \cap [form(\{w_1, w_2\})]$$

Since $\psi \wedge form(\{w_1, w_2\})$ is logically equivalent to $form(\{w_1, w_2\})$, Postulate 6 allows us to simplify the LHS:

$$[B * form(\{w_1, w_2\})] = [B * \psi] \cap \{w_1, w_2\}$$

We established earlier that $w_1 \in [B * form(\{w_1, w_2\})]$. Therefore, $w_1 \in [B * \psi] \cap \{w_1, w_2\}$, which entails $w_1 \in [B * \psi]$.

Subcase 2.2: $w_3 \in [B * \psi]$. This means $[B * \psi] \cap [form(\{w_2, w_3\})] \neq \emptyset$. Applying Postulates 7 and 8 exactly as above yields:

$$[B * form(\{w_2, w_3\})] = [B * \psi] \cap \{w_2, w_3\}$$

From our initial premise $w_2 \preceq_B w_3$ and $w_2 \notin [B]$, we know $w_2 \in [B * form(\{w_2, w_3\})]$. Therefore, $w_2 \in [B * \psi] \cap \{w_2, w_3\}$, which strictly entails $w_2 \in [B * \psi]$. Now, because $w_2 \in [B * \psi]$, the intersection $[B * \psi] \cap [form(\{w_1, w_2\})]$ is non-empty. Applying Postulates 7 and 8 to this intersection yields:

$$[B * form(\{w_1, w_2\})] = [B * \psi] \cap \{w_1, w_2\}$$

From our initial premise $w_1 \preceq_B w_2$, we know $w_1 \in [B * form(\{w_1, w_2\})]$. Thus, $w_1 \in [B * \psi] \cap \{w_1, w_2\}$, proving that $w_1 \in [B * \psi]$.

Subcase 2.3: Neither w_2 nor w_3 is in $[B * \psi]$. Since $[B * \psi]$ is a non-empty subset of $\{w_1, w_2, w_3\}$, its only remaining possible element must be w_1 . Thus, $w_1 \in [B * \psi]$.

In all possible subcases, we have proved that $w_1 \in [B * \psi]$.

Finally, since $w_1 \in \{w_1, w_3\}$, the intersection $[B * \psi] \cap [form(\{w_1, w_3\})]$ is non-empty. Applying Postulates 7 and 8 one last time yields:

$$[B * form(\{w_1, w_3\})] = [B * \psi] \cap \{w_1, w_3\}$$

Because $w_1 \in [B * \psi]$ and $w_1 \in \{w_1, w_3\}$, it follows that $w_1 \in [B * form(\{w_1, w_3\})]$.

Conclusion. Since in all subcases, $w_1 \in [B * form(\{w_1, w_3\})]$, we conclude that $w_1 \preceq_B w_3$.

3. Centeredness. We need to establish two claims:

1. If $w_1, w_2 \in [B]$, then $w_1 \preceq_B w_2$.
2. If $w_1 \in [B]$ and $w_2 \notin [B]$, then $w_1 \prec_B w_2$.

Claim (1) follows immediately from the definition above—more precisely, from $w_1, w_2 \in [B]$ it follows not only that $w_1 \preceq_B w_2$, but also that $w_2 \preceq_B w_1$.

Let us prove Claim (2). Suppose that $w_1 \in [B]$ and $w_2 \notin [B]$. From the first half of our definition of \preceq_B , it follows immediately that $w_1 \preceq_B w_2$ (since $w_1 \in [B]$).

Secondly, note that $B \not\vdash \neg form(\{w_1, w_2\})$, because the intersection $[B] \cap [form(\{w_1, w_2\})] = [B] \cap \{w_1, w_2\} = \{w_1\}$ is non-empty. Thus, the new formula is consistent with B . By the Vacuity postulate,

$$[B * form(\{w_1, w_2\})] = [B] \cap \{w_1, w_2\} = \{w_1\}$$

Because w_2 is not in this truth-set, $w_2 \notin [B * form(\{w_1, w_2\})]$. Since we also know $w_2 \notin [B]$, it fails both conditions of our definition, meaning $w_2 \not\prec_B w_1$.

Since $w_1 \preceq_B w_2$ and $w_2 \not\prec_B w_1$, it strictly follows that $w_1 \prec_B w_2$.

We have established Claim (2), and hence Centeredness.

4. Limit Assumption. Let us now prove that the **Limit Assumption** holds for \preceq_B . Note that, in a sense, the Limit Assumption is "useless" in this specific setting. It is usually assumed in order to rule out the existence of infinitely descending \prec_B -chains. However, since Φ (the set of propositional variables) is finite, \mathcal{W} is finite as well, which inherently rules out the possibility of infinitely descending chains.

Still, let us formally prove that LA holds using basic set theory. Recall that the LA reads as follows: For any sentence $\phi \in \mathcal{L}$, if $[\phi] \neq \emptyset$, then there exists at least one world $w_1 \in [\phi]$ such that $w_1 \preceq_B w_2$ for all $w_2 \in [\phi]$.

Let $\phi \in \mathcal{L}$ be such that $[\phi] \neq \emptyset$. We need to prove that there exists a minimum world in this set under the total preorder \preceq_B .

Suppose for *reductio* that no such minimum exists. That is, suppose that, for all $w \in [\phi]$, there exists a world $w' \in [\phi]$ such that $w \not\prec_B w'$. Since \preceq_B is connected, either $w \preceq_B w'$ or $w' \preceq_B w$, and so from $w \not\prec_B w'$ it follows that $w' \preceq_B w$, i.e., $w' \prec_B w$. So, the *reductio* assumption amounts to the assumption that, for all $w \in [\phi]$, there exists a world $w' \in [\phi]$ such that $w' \prec_B w$.

Let n be the total number of distinct possible worlds in the finite set $[\phi]$. Since $[\phi]$ is non-empty, we can pick some world in it, and call it $x_1 \in [\phi]$. By our *reductio* assumption, there exists $x_2 \in [\phi]$ such that $x_2 \prec_B x_1$. In turn, there exists $x_3 \in [\phi]$ such that $x_3 \prec_B x_2$, and so on.

If we repeat this process n times, we generate a strictly descending chain containing exactly $n + 1$ entries:

$$x_{n+1} \prec_B x_n \prec_B \cdots \prec_B x_3 \prec_B x_2 \prec_B x_1$$

Because our generated sequence contains $n + 1$ entries, but the set $[\phi]$ only contains n distinct worlds, it is mathematically impossible for every world in the chain to be unique. There must be some world that appears at least twice in the sequence, meaning we have $x_j = x_i$ for some step $j > i$.

By the **transitivity** of strict preference, the chain $x_j \prec_B \cdots \prec_B x_i$ strictly entails $x_j \prec_B x_i$. But since $x_j = x_i$, this means $x_i \prec_B x_i$. This is a direct contradiction, as $x_i \prec_B x_i$ means $x_i \preceq_B x_i$ and $x_i \not\prec_B x_i$, an out-and-out contradiction.

Therefore, we must reject the *reductio* assumption. There must exist at least one world $w_1 \in [\phi]$ that has no strictly more plausible world below it. Because \preceq_B is connected, this minimal world must be at least as plausible as all other worlds in the set: $w_1 \preceq_B w_2$ for all $w_2 \in [\phi]$. Thus, the Limit Assumption holds.

II. Part Two: Proving the equality $B * \phi = T(\min_B([\phi]))$.

Having established that the defined relation \preceq_B satisfies all the required structural properties (Connectedness, Transitivity, Centeredness, and the Limit Assumption), we can now prove the following equality:

$$B * \phi = T(\min_B([\phi]))$$

Because we are operating in a finite language, [Session 2 - Lemma 4](#) established that we can perfectly translate between syntax and semantics without losing information. Specifically, Property 2 of that lemma states that $T([\Gamma]) = \Gamma$ if and only if Γ is a belief set. Since the Closure postulate guarantees that the revised set $B * \phi$ is a logically closed belief set, it strictly follows that $T([B * \phi]) = B * \phi$.

Therefore, to prove our target equality, it is strictly sufficient to prove the corresponding identity:

$$[B * \phi] = \min_B([\phi]) \tag{C}$$

For if $[B * \phi] = \min_B([\phi])$ holds, $T([B * \phi]) = T(\min_B([\phi]))$ holds as well. But $T([B * \phi]) = B * \phi$ by [Session 2 - Lemma 4](#) and the fact that $B * \phi$ is a belief set by Closure. (Recall that we are assuming $*$ satisfies all AGM postulates.)

First, consider the edge case where ϕ is unsatisfiable ($[\phi] = \emptyset$). By the Success postulate, $B * \phi \vdash \phi$, meaning $[B * \phi] \subseteq [\phi] = \emptyset$. Thus, $[B * \phi] = \emptyset$. Since there are no worlds in $[\phi]$, the set of minimal worlds $\min_B([\phi])$ is trivially \emptyset as well. Thus, (C) holds.

Henceforth, assume ϕ is satisfiable ($[\phi] \neq \emptyset$). We prove the identity by mutual inclusion.

1. Prove $[B * \phi] \subseteq \min_B([\phi])$. Suppose for *reductio* that $w \in [B * \phi]$ but $w \notin \min_B([\phi])$.

By the Success postulate, $B * \phi \vdash \phi$, meaning $[B * \phi] \subseteq [\phi]$. Thus, $w \in [\phi]$. Since $w \in [\phi]$ but is not minimal in it, there must exist some world $w' \in [\phi]$ such that $w' \prec_B w$.

We evaluate this across two cases based on whether w' is a model of our original belief set B :

- **Case 1:** $w' \in [B]$. Since $w' \in [\phi]$ as well, the intersection $[B] \cap [\phi]$ is non-empty. By the Vacuity postulate, this guarantees that $[B * \phi] = [B] \cap [\phi]$. Because we assumed at the outset that $w \in [B * \phi]$, it follows that $w \in [B]$. However, by Centeredness, any two worlds in $[B]$ are equally plausible, meaning $w \preceq_B w'$. This strictly contradicts the fact that $w' \prec_B w$.
- **Case 2:** $w' \notin [B]$. Because $w' \prec_B w$, we know $w' \preceq_B w$ and $w \not\preceq_B w'$. By the definition of our ordering \preceq_B , the fact that w' is strictly preferred to w means that $[B * \text{form}(\{w, w'\})] = \{w'\}$. Since both w and w' are models of ϕ , the formula $\phi \wedge \text{form}(\{w, w'\})$ is logically equivalent to $\text{form}(\{w, w'\})$. Therefore, their truth-sets are identical. By the Superexpansion postulate (which holds unconditionally), it follows that:

$$[B * \phi] \cap [\text{form}(\{w, w'\})] \subseteq [B * (\phi \wedge \text{form}(\{w, w'\}))]$$

Which simplifies to:

$$[B * \phi] \cap \{w, w'\} \subseteq [B * \text{form}(\{w, w'\})]$$

Since we established the right side is exactly $\{w'\}$, we have:

$$[B * \phi] \cap \{w, w'\} \subseteq \{w'\}$$

Recall our starting assumption: $w \in [B * \phi]$. Since w is also obviously in $\{w, w'\}$, it must be that $w \in [B * \phi] \cap \{w, w'\}$. The subset inclusion above thus implies $w \in \{w'\}$, meaning $w = w'$. Therefore, since $w' \prec_B w$, it follows that $w \prec_B w$, which is an out-and-out contradiction.

In both cases, we reach a contradiction. Thus, we reject the *reductio* assumption. The inclusion $[B * \phi] \subseteq \min_B([\phi])$ holds.

2. Prove $\min_B([\phi]) \subseteq [B * \phi]$. Suppose for *reductio* that $w \in \min_B([\phi])$ but $w \notin [B * \phi]$. Since we assumed ϕ is satisfiable, the Consistency postulate guarantees that $B * \phi$ is consistent, meaning $[B * \phi] \neq \emptyset$. Therefore, there exists at least one world $w' \in [B * \phi]$.

Since $w' \in [B * \phi]$ and (by Success) $[B * \phi] \subseteq [\phi]$, both w and w' are models of ϕ . Thus, exactly as before, $[\phi \wedge \text{form}(\{w, w'\})] = [\text{form}(\{w, w'\})] = \{w, w'\}$. Because $w' \in [B * \phi]$ and $w' \in \{w, w'\}$, the intersection $[B * \phi] \cap [\text{form}(\{w, w'\})]$ is strictly non-empty.

Because this intersection is non-empty, it means that the revised belief set $B * \phi$ is consistent with the formula $\text{form}(\{w, w'\})$, i.e., $B * \phi \not\vdash \neg \text{form}(\{w, w'\})$. This perfectly satisfies the prerequisite for the Subexpansion postulate. Therefore, Superexpansion and Subexpansion together guarantee exact equality:

$$[B * \phi] \cap [\text{form}(\{w, w'\})] = [B * \text{form}(\{w, w'\})]$$

Recall our *reductio* assumption: $w \notin [B * \phi]$. This means w cannot be in the intersection on the left side, meaning $w \notin [B * \text{form}(\{w, w'\})]$. Since the right side must be a non-empty subset of $\{w, w'\}$, and we just established it cannot contain w , we are strictly left with:

$$[B * \text{form}(\{w, w'\})] = \{w'\}$$

Because $\{w'\}$ is the sole minimal world in this restricted set, it follows by our definition of the ordering that $w' \preceq_B w$. On the other hand, since w is minimal in $[\phi]$ (by our starting assumption) and $w' \in [\phi]$, it must be that $w \preceq_B w'$.

Since $w \preceq_B w'$ holds, but $w \notin [B * \text{form}(\{w, w'\})]$, the *only* way our definition of \preceq_B allows $w \preceq_B w'$ to be true is if its first clause is satisfied: namely, $w \in [B]$.

If $w \in [B]$, then since $w \in [\phi]$, the intersection $[B] \cap [\phi]$ is non-empty. By the Vacuity postulate, a non-empty intersection guarantees $[B * \phi] = [B] \cap [\phi]$. Since $w \in [B]$ and $w \in [\phi]$, it necessarily follows that $w \in [B * \phi]$.

This directly contradicts our *reductio* assumption that $w \notin [B * \phi]$. Thus, we reject the assumption, and the second inclusion $\min_B([\phi]) \subseteq [B * \phi]$ holds.

3. Conclusion of Part II. Having established both $[B * \phi] \subseteq \min_B([\phi])$ and $\min_B([\phi]) \subseteq [B * \phi]$, we conclude that the semantic identity (C) holds:

$$[B * \phi] = \min_B([\phi])$$

As shown at the outset of this proof, this identity guarantees our target:

$$B * \phi = T(\min_B([\phi]))$$

This completely proves the second half of the Representation Theorem. □