Single Premise Closure in the CCNCP Theory of Belief

The CCNCP theory says that S believes that *p* iff S's conditional preferences are unchanged by conditionalisation on *p*. The CCNCP stands for Conditionalisation Changes No Conditional Preferences. Formally, it looks like this

$$Bel(p) \leftrightarrow \forall A \forall B \forall q\ (A \geq_q B \leftrightarrow A \geq_{p \wedge q} B)$$

Where $Bel(p)$ means the agent believes that *p*, and $A \geq_q B$ means that the agent prefers A to B conditional on *q*. In order for the CCNCP to be a sensible theory, the three quantifiers have to be restricted, and in "Can we do without Pragmatic Encroachment?" I say a little about what the restrictions are. Here we'll be interested in three salient features of the restriction of the propositional quantifier.

- It includes all propositions that the agent is thinking about, subject to the third constraint
- It is closed under (consistent) conjunction, subject to the third constraint
- It excludes all propositions inconsistent with *p*

One of the things I left out of the earlier paper was a discussion of whether the CCNCP (which there didn't have a name) guarantees a version of single premise closure. The answer is that as it stands it doesn't, though this can be fixed with a small change. Here's the closure principle I'd like to be able to adopt.

> If the agent recognises that *p*, which she believes, entails *r*, and hence suitably adjusts all of her credences and preferences so as to be consistent with the fact that *p* entails *r*, and does not cease to believe *p* in the process of doing this, then she believes that *r*. (Moreover, though I won't argue this here, if she is justified in believing that *p* throughout, she is justified in believing that *r*.)

Here is how we might go about trying to prove this. By hypothesis, the agent in question is thinking about *r*, so it is a salient proposition. Hence for any *q* that is salient, $q \wedge r$ is salient. Hence, since the agent believes that *p*, we have the following.

(1) $\forall A \forall B \forall q\ (A \geq_{q \wedge r} B \leftrightarrow A \geq_{p \wedge q \wedge r} B)$

Since the agent recognises that *p* entails *r*, and has adjusted her preferences accordingly, her preferences conditional on $p \wedge q \wedge r$ are just the same as her preferences conditional on $p \wedge q$. So we have (2).

(2) $\forall A \forall B \forall q\ (A \geq_{p \wedge q \wedge r} B \leftrightarrow A \geq_{p \wedge q} B)$

Since the agent believes *p* we also have (3), which is just the RHS of the definition.

(3) $\forall A \forall B \forall q\ (A \geq_{p \wedge q} B \leftrightarrow A \geq_q B)$

Putting (1) to (3) together, we get (4).

(4) $\forall A \forall B \forall q\ (A \geq_{q \wedge r} B \leftrightarrow A \geq_q B)$

But that's to say that conditionalising on *r* changes no conditional preferences, so the agent believes that *r*.

The problem is that this assumes the quantifier domains are the same for *p* and for *r*, and the theory does not in general have this feature. There are two ways in which the quantifier domains can vary. First, the quantifiers over actions includes the 'actions' *believe that p*, and *don't believe that p* if, but only if, *p* is a salient proposition. (By stipulation, the agent prefers to believe that *p* iff her credence in *p* is over ½.) This won't matter to the current case – we're assuming that *p* and *r* are both salient, so those two actions are in the quantifier domain both in (1) and (4). The other problem is that the propositional quantifier also has a variable domain. It is a requirement that any proposition in that domain be consistent with the proposition that we're trying to figure out whether the agent believes. Now there are propositions that are consistent with *r* that are not consistent with *p*. For example, consider the following case where *x* and *y* are propositions, and the agent's credence function is as follows.

$$Pr(x) = Pr(y) = Pr(x \mid y) = Pr(x \mid \neg y) = 0.9$$

And assume that nothing practical turns on *x* and *y* except whether to take or decline the following bet.

Win \$50 if $\neg(x \wedge \neg y)$
Lose \$1 if $x \wedge \neg y$

Now the agent believes $x \wedge y$, for it has probability over ½ and conditionalising on it changes no conditional preferences. But she doesn't, it turns out, believe *x*, because the following is not true, where A = take the bet and B = decline the bet.

$$A \succeq_{x \wedge \neg y} B \leftrightarrow A \succeq_{\neg y} B$$

In fact the LHS is false while the RHS is true. Something has gone wrong.

The fix, I think, is to change the third clause to the following

- It excludes all propositions inconsistent either with *p* or with a belief of the agent's

Since the agent believes $\neg y$, that can't be in the quantifier domain when we're trying to work out whether the agent believes that *x*. That makes some amount of sense. We don't want to say that the agent fails to believe that *p* because learning that *p* for sure changes the agent's preferences conditional on something they believe to be false. A state of confidence in *p* doesn't amount to belief if conditionalising on *p* changes preferences over live options, but preferences conditional on something the agent believes to be false are not exactly live in any interesting sense. So this is a bug, to be sure, but one that can be easily patched.