(editorial comments not included)

From Bart:

-> In general, there's a problem that the paper is obviously written by different people, and different notation. Compare for instance "Adv^{QuEME}_prp" of the classical part with "prp-Adv_{QuEME}" in the quantum part. Also the writing style differs significantly over the paper. [NOTATION DONE, REST TO CHECK]

-> Section 3.2: I discussed this with Ritam already, but basically you conjecture that it is possible to tighten a mirror theory proof that doesn't even exist yet. I think this requires more justification. Also, this shouldn't belong in "Summary of results" but a separate section perhaps. [DONE]

-> The structure is reasonably okay but sometimes it is unclear what you are doing. There are many things, like above point but many other things that just get justified well enough. For example, you present the general scheme and then a quantum attack on a specific variant. In particular, the quantum attack is for a SPECIFIC mixing function. Is it possible to extend this quantum attack to eliminate any mixing function that combines the two inputs before they are fed to the primitive. E.g., suppose you build
M(hatL,hatR) = linear'(f(linear(hatL,hatR)),hatL,hatR)
for any n-to-n-bit (!) function f, can your attack still work? That would be a very good justification for using a cryptographic scheme of 2n-to-n bits. Now, your mixing function QuEME comes a bit out of the blue given that the attack of Section 4.1 only attacks the Lai-Massey mixing. [Done]

-> The attack of Section 4.1 itself is too technical and compressed. When Paul explained it yesterday, I got it, but I didn't get it when reading the paper and now I somewhat get it. This one needs more clarification (particularly, if it would be extended to any n-bit cryptographic primitive). [Done]

-> Algorithm 1: also talks about S(...) which is (?) undefined [Done]

Proof of Section 5. We (I and Ritam) may disagree but to me it is weird to use standard model for E1,E2,E3,E4 and ideal model for E, but still give no access to the keyless ideal cipher. Solutions:
- put an extra key and add it to X before it goes into the scheme. (perhaps not needed as option 2 works as well).
- standard-model step like in Minematsu 2009's TBC paper
- ideal cipher model proof all the way (adds a term 4q/2^n + q^2/2^n). [DONE]

-> Step 2 of the sampler: unclear whether this is non-empty. Later on, you'll use the mirror theory for that. Best to remark it. [DONE]

-> Proof used many undefined notation (what is a path, what is Pr_0, ...) Of course I know what it is, but it should be clarified. [DONE]

-> Same for other parts of the paper, by the way. [TO CHECK]

-> Section 6, I really like the idea that you did this and what you did, but it deserves more time and words. In particular, the stuff starting

from Approximation consists of loose sentences without a fluent storyline. Also, the outcome is not clear. We already discussed this, but the n was unclear. Implementation aspects? Explanation about the steps in algorithm 2 are too drafty. How far are you away of the actual bound of Patarin? This last question is particularly interesting I think. [PARTLY REWRITTEN, TO CHECK/IMPROVE]

-> Quantum proof:
- Proposition 4 talks about random subsets L,R, but L and R are values in QuEME right?
- Game 3: shouldn't S_n(r) be D_n(r)? Also, mentioned before, completely different structure.
- Corollary 1: you assume that the classical B-advantage is O(...). Why do you assume that? And not just prove that? Can't you relate it to the proof of Sections 5? Also, weird to put an assumption in a corollary. [TO DO]

-> In the instantiation of Section 9:
- of course it makes sense to do keys like this, and I think it's nicely sold.
- also here, notation odd: is it k_i or K_i?
- for the classical proof, basically you can rely on a really simple type of related-key security to make the step to pi1,...,pi4. Of course, for the quantum proof that would be much harder.
- what happens in the classical proof if you take pi3=pi1 and pi4=pi2? The proof would become harder but it would be a much cleaner and easier to sell variant. Of course, again, for the quantum proof that would be annoying. [TO CHECK]

From Reviewer A:

-> It seems that the BHT part of Algorithm 1 cannot directly obtain the entangled state of R_0 and R_1 satisfying this condition, even though without measurement. Because that the BHT is essentially an amplitude amplification algorithm, which eventually produces a uniform superposition of states colliding with elements in the target set. [DONE]

-> In the Grover-meet-Simon part of Algorithm 1, please illustrate the specific form of the function with hidden Boolean period? Whether the function can be accessed by the adversary? [DONE]

-> Is Algorithm 2 an algorithm proposed in the previous literature or a new algorithm proposed in this work? [TO CHECK]

-> The correctness of algorithm 1 is doubtful. Besides, the quantum security proof method is too rough, and the result security bound is not tight enough. [TO IMPROVE WRITEUP]

From Reviewer B:

-> As far as I see, all the proofs of the provable security results, both classical and quantum, have a critical error. In the middle part, the

mixing layer, of QuEME is the ideal-cipher, a keyless primitive (line -2 in Sect. 5 and also in the instantiation in Sect. 9). In the very first step of the proof (in Sect. 5.1), this is replaced with a tweakable random permutation, a keyed primitive. This does not work. In the former, the adversary should be given oracle access to E, and this power suddenly disappears after the replacement. This invalidates all the provable security results of the paper, and this is particularly relevant to the Q2 adversary that can query all the input-key (or output-key) to the ideal-cipher with one quantum query. The paper has to be revised from the security definitions incorporating the ideal-cipher, or the scheme should be re-designed. [DONE]

-> Section 9 is the weakest part of the paper. No cryptanalytic attempt has been made to evaluate its security. The authors may want to present their best attack on the reduced round version, say 6-round version, to justify their choice of the number of rounds. The fact that the instantiation has only 2n-bit keys is not justified. [TO DO]

-> Page 3. Please add a figure describing EME. The description in Sect. 2.1 is painful to parse. [TO DO]

-> Section 3.4 is poorly presented, and this section requires a major rewrite. [TO DO]

-> Section 4.1 is very poorly presented. The authors should add the description of Algorithm 1 as texts in the main body. [TO DO]

From Reviewer C:

-> Can you clarify the model in your security proofs? What kind of access to E_i (\pi_i) is permitted? [DONE]

-> The quantum attack attempts to recover the key of one of the n-bit block ciphers. It adds an additional Claw Finding layer on an existing technique, which combines Grover search with Simon's algorithm. This new approach makes sense, but there is no proof to verify it. Some ingredient of the algorithm also needs to be clarified. For instance, the algorithm assumes BHT creates uniform superposition on Claws. This property should be justified and stated in your description of BHT algorithm. [Done]

-> Algorithm 1, the quantum attack, is difficult to parse. It is helpful to present it in a modular way as much as possible. [Done]

-> Section 4.1., make a formal theorem statement about the superposition attack (and give its analysis). Incidentally, define k. [Done]

-> In Theorem 4, state clearly which conjecture from mirror theory is used. Now it only refers to section 3.2. In the proof, when clearly state it when these conjectures are used. Also give pointers to the CCA proof in the appendix. [TO DO]

-> Why a separate CPA security proof in A.2. once CCA-security has been established? [TO JUSTIFY]

-> I also suggest calling Theorem 1 and 2 Conjectures to avoid confusion.
[TO DISCUSS]