

# Fairness in the Wild: Secure Atomic Swap with External Incentives

Hao Chung\*  
Carnegie Mellon University  
chunghaoqc@gmail.com

Elisaweta Masserova†  
Carnegie Mellon University  
elisawem@andrew.cmu.edu

Elaine Shi\*  
Carnegie Mellon University, Oblivious Labs Inc., and 0xPARC  
elainersh@gmail.com

Sri AravindaKrishnan Thyagarajan\*  
University of Sydney  
aravind.thyagarajan@sydney.edu.au

## Abstract

Atomic swaps enable asset exchanges across blockchains without relying on trusted intermediaries, and are a key component of decentralized finance (DeFi) ecosystems. Recently, Chung, Masserova, Shi, and Thyagarajan introduced *Rapidash* (Financial Cryptography 2025), an atomic swap protocol that remains incentive compatible under user-miner collusion, by ensuring that the honest strategy forms a coalition-resistant Nash equilibrium. However, their model assumes a closed system where players act solely based on internal protocol incentives. In practice, participants may be influenced by *external incentives* such as off-chain rewards or adversarial bribes, which can undermine such equilibrium guarantees.

In this work, we introduce a new game-theoretic notion, *bounded maximin fairness*, which ensures that honest participants remain protected against rational adversaries with arbitrary but bounded external incentives. We construct an atomic swap protocol that satisfies this notion, while preserving the equilibrium properties of prior work in the absence of external influence. As we show, our protocol is easy to implement and can be instantiated even in Bitcoin’s limited scripting language.

---

\*Supported by NSF awards 2212746, 2044679, 1704788, a Packard Fellowship, a generous gift from the late Nikolai Mushegian, a gift from Google, and an ACE center grant from Algorand Foundation.

†Supported by NSF Grant 1801369, the CONIX Research Center, the Defense Advanced Research Projects Agency under contract FA8750-17-1-0059, a gift from Bosch, and the CBI postdoctoral fellowship.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Our Contributions . . . . .	4
1.2	Related Work . . . . .	4
<b>2</b>	<b>Technical Overview</b>	<b>5</b>
2.1	Knowledge-Coin Exchange . . . . .	5
2.2	From Knowledge-Coin Exchange to Atomic Swap . . . . .	8
<b>3</b>	<b>Model</b>	<b>10</b>
3.1	Blockchain and Smart Contracts . . . . .	10
3.2	Players and Strategy Spaces . . . . .	11
3.3	Game-Theoretic Properties . . . . .	12
3.3.1	New Definition: Bounded Maximin Fairness . . . . .	12
3.4	Atomic Swap . . . . .	14
<b>4</b>	<b>Secure Atomic Swap</b>	<b>15</b>
<b>5</b>	<b>Proof of Theorem 4.1</b>	<b>21</b>
5.1	Achieving CSP Fairness . . . . .	21
5.2	Achieving Bounded Maximin Fairness . . . . .	25
5.2.1	Irrational Strategies . . . . .	25
5.2.2	Against Externally Incentivized Bob-Miner Coalition . . . . .	27
5.2.3	Against Externally Incentivized Alice-Miner Coalition . . . . .	32
5.2.4	Combine Everything Together . . . . .	36
5.3	Achieving Dropout Resilience . . . . .	38
<b>6</b>	<b>Rational Defection from Grand Coalition</b>	<b>38</b>
6.1	Disincentivizing the Grand Coalition Absent of External Incentive. . . . .	38
6.2	Disincentivizing the Grand Coalition in the Presence of External Incentive. . . . .	39
<b>7</b>	<b>Instantiation</b>	<b>40</b>
7.1	Ethereum Instantiation . . . . .	40
7.2	Bitcoin Instantiation . . . . .	43
7.2.1	Instantiating $\text{CONTRACT}^B$ with Bounded Maximin Fairness . . . . .	43
7.2.2	Instantiating $\text{CONTRACT}^A$ with Bounded Maximin Fairness . . . . .	44

# 1 Introduction

In distributed systems and cryptographic protocols, fairness is a central yet elusive goal. However, achieving the strongest form of fairness — either all parties receive the output or none do — is often impossible for general multi-party functionalities without heavy assumptions such as honest majority or trusted execution environments [PG99]. This limitation has motivated the search for relaxed fairness notions that still provide meaningful guarantees for honest participants, even when full fairness is unattainable.

One promising direction is to embrace a game-theoretic perspective, recognizing that real-world participants are often *rational* parties, motivated by self-interest rather than classified as either honest or malicious in the classical cryptographic sense. Protocols in this paradigm aim to ensure that no coalition of rational parties has an incentive to deviate from the prescribed behavior. This shift in perspective has already led to substantial progress in domains previously blocked by impossibility results. For instance, in the case of fair coin tossing recent game-theoretic protocols [WAS22, TSW24, CGL<sup>+</sup>18] obtained fairness even under dishonest majority — a setting where classical cryptographic notions of fairness have long been known to be impossible to achieve.

We revisit another fundamental primitive in decentralized environments: atomic swaps. These protocols enable two mutually distrustful parties to exchange digital assets across different chains without relying on a trusted intermediary. However, the blockchain setting introduces new challenges: miners can selectively include, exclude, or reorder transactions, and users may collude with the miners to extract unfair advantage. To address this, a recent work dubbed Rapidash [CMST22b] laid a foundational game-theoretic framework for atomic swaps that remain incentive-compatible even in the presence of user-miner collusion. The Rapidash protocol was proven to satisfy cooperative strategy proofness (CSP-fairness), which guarantees that the coalition’s utility is maximized when all players in the coalition follow the honest strategy. Thus, informally, in Rapidash’s atomic swap the honest strategy formed a coalition-resistant Nash equilibrium.

Yet, Rapidash left one important question open: what if parties have access to external incentives, such as side payments or off-chain contracts with external entities? In realistic settings, particularly in decentralized finance (DeFi), participants do not operate in isolation. They may be subject to off-protocol pressures or rewards that alter their incentives. Consider an atomic swap where one party, Alice, is a major liquidity provider on a decentralized exchange, and her ability to maintain this role depends on successfully completing a swap with Bob, e.g., acquiring BTC to rebalance her position. Suppose a third party, Eve, stands to gain market dominance if Alice’s trading capacity is disrupted. In this scenario, Eve may have a strong incentive to bribe Bob to misbehave during the atomic swap — ensuring that Alice does not receive the BTC in time, thereby impairing her operations. Note that in this case, the CSP guarantee is insufficient, as it offers no protection when Bob is willing to adopt a strategy that does not maximize his utility within the swap itself. Without external incentives, such scenarios are not of concern, since it is reasonable to assume that Bob would always aim to maximize his own profit within the atomic swap. However, in the presence of off-protocol incentives, Bob may find it optimal to even intentionally *lose money* within the atomic swap protocol as long as Eve’s compensation offsets this loss.

As out-of-band incentives clearly violate the closed-world assumption underlying many rational protocol models and can undermine the guarantees of even game-theoretically secure constructions, in this work we ask ourselves the following question:

*Is it possible to provide an atomic swap protocol that is secure against strategic players with external incentives?*

## 1.1 Our Contributions

We answer this question positively, and our contribution here is two-fold. First, we formalize a new game-theoretic notion, called *bounded maximin fairness*, which captures the idea that honest players should not be harmed by the presence of external incentives, and it may be of independent interest in a broader context. Second, we design an atomic swap protocol that satisfies this notion. Moreover, we present the Bitcoin instantiation of our protocol showcasing the practicality of our solution.

**New definition: bounded maximin fairness.** As explained above, the notion of CSP fairness does not capture scenarios where parties can be influenced by external incentives. Intuitively, when such incentives are unbounded, they can motivate the participating parties to behave arbitrarily maliciously. In this sense, the strongest possible game-theoretic notion one can hope for is that even when the strategic coalition can be arbitrarily malicious, honest players should not be harmed. This idea underpins *maximin fairness*, a notion introduced in recent works [PS17, CGL<sup>+</sup>18, WAS22]. Maximin fairness implicitly assumes that the external incentives may be arbitrary and unbounded, which explains why the strategic players’ behavior may appear arbitrarily malicious from the perspective of the present protocol. Maximin fairness, however, is a very stringent requirement, and insisting on such a strong notion of fairness may severely constrain the design space or even lead to impossibility results in some applications.

For atomic swap, we currently do not know how to achieve maximin fairness given strategic miners. Instead, we suggest a meaningful relaxation called *bounded maximin fairness*. Unlike its more stringent counterpart, bounded maximin fairness allows external incentives to be arbitrarily large but assumes *there is an upper bound on the external incentives and the bound is known to the protocol*. We want to guarantee that honest participation will not lead to negative utility as long as the other strategic players behave rationally in light of the arbitrary but bounded external incentives. This guarantee nicely complements equilibrium-type guarantees such as CSP fairness — essentially, bounded maximin provides a safety net for honest parties if they suspect that there may be external influence at play. In other words, simple-minded, non-strategic players can always feel safe participating in the protocol.

**Secure atomic swap.** We design an atomic swap protocol that satisfies both bounded maximin fairness and CSP fairness. Additionally, our protocol has *dropout resilience*, which ensures that an honest player is protected from loss if the other player plays honestly, but drops offline at any point during the protocol. This property is important in practice since the honest player may go offline due to loss of password or network outage.

Finally, we implement our protocol for blockchains that support general smart contracts, and also demonstrate how it can be instantiated with Bitcoin scripts — showcasing its practicality even on more restrictive platforms like Bitcoin.

## 1.2 Related Work

Over the past decades, significant effort has been devoted to achieving fair exchange on blockchains. Numerous elegant solutions have been proposed under the assumption that miners behave honestly. For example, Choudhuri et al. [CGJ<sup>+</sup>17] presented a solution for achieving fair exchange and fair multi-party computation using a public bulletin board and extractable witness encryption. For atomic swaps, several approaches have been developed [Her18, MMS<sup>+</sup>, vdM19, MD19, ZDBN19]. Typically, atomic swaps are based on hashed time-lock contracts (HTLCs) [Her18] or a novel cryptographic primitive known as *adaptor signatures* [DH20].

Another line of work focuses on achieving fair exchange when miners may behave strategically [CMST22b, WSZN22, CMST22a]. A series of papers [WHF19, Bon16, MHM18] demonstrated how users can bribe miners to censor competing transactions. To improve the robustness of HTLCs under such adversarial conditions, Tsabary et al. [TYME21] proposed *MAD-HTLC*, which requires the initiating party to lock up extra collateral. Although MAD-HTLC mitigates certain bribery attacks, it fails to achieve CSP-fairness in the presence of general user-miner collusion. More recent works, such as Rapardash [CMST22b] and He-HTLC [WSZN22], extend the ideas behind MAD-HTLC. These works propose new *knowledge-coin exchange* protocols, which achieve the same functionality as HTLCs, while providing provable CSP-fairness guarantees. However, none of these protocols satisfies the bounded maximin fairness.

Interestingly, as pointed out by Chung et al. [CMST22b], even if the underlying knowledge-coin exchange protocols are CSP-fair, the atomic swap protocol built upon them may not inherit this property. This highlights that game-theoretic properties like CSP-fairness are not necessarily composable. Therefore, a carefully designed atomic swap protocol is required to ensure CSP-fairness. A more detailed discussion on achieving CSP-fairness in atomic swaps is provided in Section 2.2.

Beyond fair exchange, a notable category of external incentives in blockchain systems is *oracle manipulation* attacks. An oracle is a service that provides real-world data to smart contracts on a blockchain. Some oracles are themselves implemented as smart contracts, where users stake tokens and receive rewards for supplying data. When the outcomes of other applications, e.g., prediction markets, depend on the oracle’s data, these applications introduce external incentives that may motivate oracle participants to submit incorrect information in order to manipulate outcomes. It was reported that over \$400 million was lost to oracle manipulation attacks in 2022 [ora22].

## 2 Technical Overview

The goal of an atomic swap is to enable two parties to exchange their cryptocurrencies across two different blockchains. Before diving into our solution, we first introduce a helpful simpler primitive, *knowledge-coin exchange*. In Section 2.1, we review the existing knowledge-coin exchange protocols, and explain why they fail to achieve bounded maximin fairness even with limited external incentives. Then, we explain the necessary modifications to achieve bounded maximin fairness. In Section 2.2, we explain how to realize atomic swap based on knowledge-coin exchange. We highlight the key challenges of designing an atomic swap protocol, and summarize the key ideas of our construction. In the following, we use “strategic players,” or “coalition,” to describe colluding parties (that can potentially exhibit malicious behavior).

### 2.1 Knowledge-Coin Exchange

**Hash timelock contract (HTLC).** In a knowledge-coin exchange, Bob wants to buy a secret  $s$  from Alice at a price of  $\$v$  coins. In practice, such functionality is often realized by *hash timelock contracts* (HTLCs), which work as follows. Alice first sends  $h_s = H(s)$  to Bob, where  $H(\cdot)$  denotes a cryptographic hash function. Then, Bob deposits  $\$v$  into the contract upfront, and the contract is parametrized with the hash value  $h_s$  and a timeout  $T$ .

#### HTLC

- On receiving  $s$  from Alice such that  $H(s) = h_s$ , send  $\$v$  to Alice.

- After  $T$ , on receiving `ping` from Bob, send  $\$v$  to Bob.

Above, `ping` is a predetermined message which triggers the corresponding *activation point*, i.e., the corresponding part of the smart contract logic. The first activation point allows Alice and Bob to exchange the secret and the payment, and the second one allows Bob to recover its deposit if Alice drops offline. Since the total deposit is only  $\$v$ , the two activation points are mutually exclusive, i.e., only one of them can be activated and only once.

While HTLC is dropout resilient, it is unfortunately not CSP fair in the presence of user-miner coalitions. If Alice offers a transaction fee of  $\$f$ , then as long as Bob bribes each miner  $\$f + \$\epsilon$  (for some small  $\epsilon > 0$ ) for excluding Alice’s transaction, rational miners will take the bribe [TYME21]. This bribery attack makes sense for Bob as long as  $\$v > \$f \cdot T$ . It may seem like HTLC is secure as long as  $T \cdot \$f$  is sufficiently large. However, Tsabary et al. [TYME21] showed new attacks where the cost to Bob is not dependent on  $T$ .

**Prior work: Achieving CSP-fairness.** To address user-miner collusion, Chung et al. [CMST22b] proposed a new protocol that achieves CSP-fairness. In this protocol, Alice first computes  $h_s = H(s)$ . Additionally, Bob samples a random string  $pre_b$ , and computes  $h_B = H(pre_b)$ . They deploy the following smart contract parametrized by hash values  $h_s$  and  $h_b$ , the timeout  $T_1$  and  $T_2$ , the value of the secret  $\$v$ , the collateral  $\$c_b$ , and a predetermined small amount  $\$e$ .

#### CSP-fair knowledge-coin exchange contract

*/\* Params:  $(h_s, h_b, T_1, T_2, \$v, \$c_b, \$e)$ , Bob deposits  $\$v + \$c_b$ . \*/*

$P_{\text{default}}$ : On receiving  $z$  from Alice s.t.  $H(z) = h_s$ , send  $\$v$  to Alice and  $\$c_b$  to Bob.

$P_{\text{refund}}$ : Time  $T_1$  or greater: on receiving  $z$  from Bob s.t.  $H(z) = h_b$ , do nothing.

$C_{\text{refund}}$ : At least  $T_2$  after  $P_{\text{refund}}$  is activated: on receiving `ping` from anyone, send  $\$v + \$c_b$  to Bob.

$C_{\text{burn}}$ : On receiving  $(z_1, z_2)$  from any  $P$  s.t.  $H(z_1) = h_s$  and  $H(z_2) = h_b$ , send  $\$e$  to player  $P$ .  
All remaining coins are burnt.

Figure 1: Rapidash’s CSP-fair knowledge-coin exchange contract [CMST22b].

In the above construction, the activation points starting with the same capital letter are mutually exclusive. For example, once  $P_{\text{default}}$  is activated,  $P_{\text{refund}}$  cannot be activated. We adopt the same convention in this work. The formal specification of writing a smart contract is given in Section 3.1.

To complete the knowledge-coin exchange, Bob deposits the intended payment  $\$v$ , and an additional  $\$c_b$  amount of collateral into the contract. Then, Alice sends  $s$  to  $P_{\text{default}}$  as soon as Bob’s deposit takes effect. In this case, Alice receives the payment  $\$v$ , and Bob learns  $s$  and gets his collateral back all at once. If Alice fails to post  $s$  by time  $T_1$ , Bob sends  $pre_b$  to  $P_{\text{refund}}$  at time  $T_1$  to request a refund. The activation point  $P_{\text{refund}}$  merely allows Bob to express his intent to request a refund. The actual refund happens when  $T_2$  time has passed since the activation of  $P_{\text{refund}}$  — at this point, Bob sends to  $C_{\text{refund}}$  which actually sends him the refund. An honest miner always includes all outstanding transactions in any block it mines. Importantly, if an honest miner has observed both  $s$  and  $pre_b$  contained in the transactions posted, it will immediately post  $(s, pre_b)$  to  $C_{\text{burn}}$ , and this transaction will always be ordered in front of others in the block it mines.

The knowledge-coin protocol above achieves CSP-fairness. The key insight here is that the activation point  $C_{\text{burn}}$  serves as a “bomb”. Suppose that the honest Alice has posted  $s$ . Now, should a strategic Bob-miner coalition post  $pre_b$  to  $P_{\text{refund}}$  in an attempt to get refunded and thus

get the secret  $s$  for free, both  $s$  and  $pre_b$  will be publicly known. Note that the coalition has to wait at least  $T_2$  amount of time before the actual refund  $C_{\text{refund}}$  can be activated. During this  $T_2$  window, if any non-colluding miner mines a block, it will trigger the bomb by posting  $(s, pre_b)$  to  $C_{\text{burn}}$ , which will cause Bob to lose its collateral. Thus, if a Bob-miner coalition wishes to get the secret for free, it has to take a gamble that it will be able to mine all blocks in the  $T_2$  window.

**Why Rapidash does not satisfy bounded maximin fairness.** Unfortunately, the above protocol does not satisfy bounded maximin fairness, as an externally motivated Alice can harm Bob. Indeed, Alice can withhold  $s$  and wait until Bob sends  $pre_b$  to  $P_{\text{refund}}$ . Then, Alice can send  $s$  and  $pre_b$  to  $C_{\text{burn}}$  to redeem  $\$e$ , thus making Bob lose his collateral  $\$c_b$ . As long as Alice is compensated externally by  $\$v - \$e$ , she is incentivized to trigger  $C_{\text{burn}}$  instead of sending  $s$  to  $P_{\text{default}}$  to earn  $\$v$ .

**Bounded maximin fairness: definition subtleties.** Before we explain our modifications to the above protocol towards one that is better aligned with our bounded maximin fairness notion, we first discuss the proposed fairness definition itself. At a first glance, a natural way to define bounded maximin fairness is to require that the *optimal strategy* that maximizes the externally incentivized coalition’s utility does not harm any honest player. It implicitly assumes that the coalition *knows* the optimal strategy. However, the external incentive can be an arbitrary function of the blockchain states, and finding the optimal strategy may not be computationally efficient. To capture the security that defends against probabilistic polynomial-time (PPT) players, we define a set of *blatantly irrational strategies*  $\bar{\mathcal{R}}$ , where a strategy  $S$  is in  $\bar{\mathcal{R}}$ , i.e., it is blatantly irrational, if one can efficiently find another strategy  $S'$  such that the coalition’s utility when adopting  $S'$  is strictly larger than that of  $S$ . Then, we say that a protocol is bounded maximin fair if honest players’ utility is non-negative when the coalition adopts any strategy outside  $\bar{\mathcal{R}}$ . We provide the formal definition in Section 3.3.1.

**Our knowledge-coin exchange.** The challenge in designing a bounded-maximin-fair protocol is to make sure *every strategy that may harm the honest players is blatantly irrational*. We achieve this by ensuring that the following conditions hold:

1. Harming the honest player necessarily puts the adversary at risk of triggering the bomb.
2. The utility of the externally incentivized coalition strictly decreases if the bomb activation point gets triggered.
3. The last message that contributes to the bomb getting triggered is always being sent by a strategic player.

The last two conditions ensure that activating the bomb is a blatantly irrational strategy: Whenever a strategic player is about to send a message which would trigger a bomb, their utility is strictly improved if they instead simply stop sending messages from that moment on. Therefore, given any strategy  $S$ , we only need to check whether stopping to send messages is a better strategy than continuing with  $S$ , which is efficiently verifiable.

To ensure that the second condition holds, we set the collateral high enough to ensure that the external incentive cannot offset the loss which occurs when a bomb is triggered. As we know that an externally incentivized Alice can harm Bob, we have Alice supply the collateral as well (and burn it together with Bob’s when a bomb is triggered).

At a first glance, the second condition may seem trivial, because how and why would honest players trigger the bomb? However, surprisingly, the CSP-fair contract specified in Figure 1 does not satisfy this condition. Consider the following: A strategic Bob posts  $pre_b$  publicly even “before

making his deposit”. Then, once Bob made his deposit, the honest Alice follows the protocol and posts  $s$  to  $P_{\text{default}}$ . At this moment, both  $s$  and  $pre_b$  are publicly known, so the miner can post  $(s, pre_b)$  to  $C_{\text{burn}}$  to trigger the bomb.

To prevent the scenario above, we regulate the timing when the strategic Bob would want to reveal  $pre_b$ . To do so, we add the following activation points, where  $h_b = H(pre_b)$ .

$B_{\text{defuse}}$ : On receiving `ping` from anyone, do nothing.

$B_{\text{burn}}$ : On receiving  $z$  from anyone  $P$  such that  $H(z) = h_b$ , send  $\$e$  to player  $P$  and burn remaining coins.

The effect of  $B_{\text{burn}}$  is similar to  $C_{\text{burn}}$  as it burns most of the deposits, but  $B_{\text{burn}}$  can be triggered simply by  $pre_b$ . Recall that the activation points of the same type are mutually exclusive. Once  $B_{\text{defuse}}$  has been activated,  $B_{\text{burn}}$  cannot be activated. Thus, the purpose of  $B_{\text{defuse}}$  is to “defuse” the bomb of  $B_{\text{burn}}$ , i.e., a strategic Bob would only want to post  $pre_b$  after  $B_{\text{defuse}}$  has been activated. Since  $B_{\text{defuse}}$  can only be triggered when the deposits have been made, this ensures that a strategic Bob would not want to post  $pre_b$  before making his deposit (which was the problem in the scenario considered above). Intuitively, posting  $pre_b$  before the deposits is now a blatantly irrational strategy. Consequently, the resulting contract is as follows.

#### Our knowledge-coin exchange contract

*/\* Params:  $(h_s, h_b, T_1, T_2, \$v, \$c_b, \$e)$ , Bob deposits  $\$v + \$c_b$ , Alice deposits  $\$c_a$ . \*/*

$B_{\text{defuse}}$ : On receiving `ping` from anyone, do nothing.

$B_{\text{burn}}$ : On receiving  $z$  from anyone  $P$  such that  $H(z) = h_b$ , send  $\$e$  to player  $P$  and burn remaining coins.

$P_{\text{default}}$ : On receiving  $z$  from Alice s.t.  $H(z) = h_s$ , send  $\$v + \$c_a$  to Alice and  $\$c_b$  to Bob.

$P_{\text{refund}}$ : Time  $T_1$  or greater: on receiving  $z$  from Bob s.t.  $H(z) = h_b$ , do nothing.

$C_{\text{refund}}$ : At least  $T_2$  after  $P_{\text{refund}}$  is activated: on receiving `ping` from anyone, send  $\$c_a$  to Alice and send  $\$v + \$c_b$  to Bob.

$C_{\text{burn}}$ : On receiving  $(z_1, z_2)$  from any  $P$  s.t.  $H(z_1) = h_s$  and  $H(z_2) = h_b$ , send  $\$e$  to player  $P$ . All remaining coins are burnt.

## 2.2 From Knowledge-Coin Exchange to Atomic Swap

Intuitively, an atomic swap can be realized by two knowledge-coin exchange instances, one on each blockchain. Suppose Bob has  $x_b$  coins on BobChain (denoted  $\mathbb{B}x_b$ ), and Alice has  $x_a$  coins on AliceChain (denoted  $\mathbb{A}x_a$ ). The goal is for Bob to exchange his  $\mathbb{B}x_b$  for Alice’s  $\mathbb{A}x_a$ . The blueprint for atomic swaps via knowledge-coin exchange composition is as follows: Alice samples a random string  $pre_s$ , and sends  $h_s = H(pre_s)$  to Bob. Then, she deposits a prescribed amount into a knowledge-coin exchange on AliceChain parametrized with the hash value  $h_s$ . Similarly, Bob deposits  $\mathbb{B}x_b$  into another knowledge-coin exchange on BobChain parametrized with the same hash value  $h_s$ . Once both Alice and Bob have deposited their coins, Alice sends  $pre_s$  to the contract on BobChain to obtain  $\mathbb{B}x_b$ . At this moment,  $pre_s$  is publicly known, and Bob can send  $pre_s$  to the contract on AliceChain to obtain  $\mathbb{A}x_a$ .

**Naive composition.** We can hope that this logic can be used to satisfy bounded maximin and CSP fairness, if we use our knowledge-coin exchange scheme from above as a base. Towards

this, in addition to  $pre_s$ , Alice samples a random string  $pre_a$  to help with the refund procedure. Similarly, Bob samples a random  $pre_b$ . They deploy a contract parametrized with the hash values  $h_s, h_a = H(pre_a)$  on AliceChain, and another contract parametrized with  $h_s, h_b = H(pre_b)$  on BobChain. They also put the prescribed deposits and collaterals into the contracts. Ideally, when everyone is honest, Alice will send  $pre_s$  to  $P_{\text{default}}^B$  on BobChain (henceforth, the superscript B indicates that the activation point is in the contract on BobChain), and Bob uses the same  $pre_s$  to activate  $P_{\text{default}}^A$  on AliceChain (similarly, the superscript A indicates that  $P_{\text{default}}^A$  is on AliceChain) to complete the atomic swap. If the other party drops out, Alice and Bob can use  $pre_a$  and  $pre_b$  to get refunded, respectively.

However, as Chung et al. [CMST22b] pointed out, the direct composition of two CSP-fair knowledge-coin exchanges does not yield a CSP-fair atomic swap. Here we briefly review the issues and their solution, as it is crucial to our construction. The issue is that Alice-miner coalition can withhold  $pre_s$ , and try to get refunded on AliceChain. Of course, at this moment, the honest Bob will also try to get refunded on BobChain by sending  $pre_b$ , but Alice-miner coalition defers Bob's refunding transaction. After Alice successfully gets refunded on AliceChain, she immediately sends  $pre_s$  to  $P_{\text{default}}^B$  to obtain  $\mathbb{X}_b$ . Unlike the context of knowledge-coin exchange where Bob wants to buy the secret  $pre_s$  at certain price, in atomic swap,  $pre_s$  is just a random string facilitating the exchange. Once Alice gets refunded on AliceChain,  $pre_s$  is worth nothing to Bob. Consequently, Alice-miner coalition can get  $\mathbb{X}_b$  for free!

To address this issue, we allow the bomb  $C_{\text{burn}}^A$  on AliceChain to be triggered by  $pre_a$  and  $pre_b$ . Consequently, if Alice-miner coalition withholds  $pre_s$  so that the honest Bob tries to get refunded by  $pre_b$ , Alice-miner coalition is disincentivized to send  $pre_a$ . This patch indeed fixes the attack above, while we lose the dropout resilience for Alice. If Alice's deposit on AliceChain is delayed, e.g., due to a network congestion, Bob will send  $pre_b$  to BobChain to get refunded. At this moment, after the deposit is finalized, the honest Alice could not get refunded by  $pre_a$  because the bomb  $C_{\text{burn}}^A$  can be triggered by  $(pre_a, pre_b)$ .

A key challenge of designing an atomic swap protocol is to find the right balance for how easy the player can be refunded. If it is too easy, it may become risk-free to attack the honest player; if it is too difficult, the honest player's coins may be locked simply due to network congestion.

**Two-phase preparation.** Chung et al. [CMST22b] proposed a two-phase preparation to address the above issue. They locked  $P_{\text{default}}^B$  and  $C_{\text{burn}}^B$  on BobChain with an additional hash  $h_c$  of a random string  $pre_c$  sampled by Bob. Bob publishes  $pre_c$  if the deposits into the contracts on both blockchains are finalized in a timely manner. The extra lock  $h_c$  then helped to distinguish the two cases depending on whether Alice strategically withholds  $pre_s$  or not.

1. If the deposit transactions are not finalized in time, Bob will never publish  $pre_c$ . In this case, Bob will send ping to  $P_{\text{refund}}^A$  and send  $pre_b$  to  $P_{\text{default}}^B$  to get refunded. Because  $P_{\text{default}}^B$  is locked with  $h_c$ , Alice cannot cash out  $\mathbb{X}_b$  from  $P_{\text{default}}^B$ , and Bob can safely help Alice get refunded.
2. If the deposit transactions are finalized in time, then Bob will publish  $pre_c$ . Here, if Bob ever needs to post  $pre_b$  to  $P_{\text{refund}}^B$ , it must be because Alice-miner coalition withholds  $pre_s$ .

**Bounded-maximin-fair atomic swap.** We adapt the ideas above and combine these with our knowledge-coin exchange to achieve bounded-maximin-fair atomic swap. Specifically, Alice and Bob deploy a knowledge-coin exchange contract on AliceChain with the following modifications:

- $C_{\text{burn}}^A$  can be triggered by  $pre_a$  and  $pre_b$ . This is to prevent Alice-miner coalition from withholding  $pre_s$  on BobChain.

- Add an activation point  $A_{\text{burn}}^A$ :
  - On receiving  $z_1$  from anyone  $P$  such that  $H(z_1) = h_a$ , or on receiving  $z_2$  from anyone  $P$  such that  $H(z_2) = h_b$ , send  $\mathbb{A}\epsilon^A$  to player  $P$ .
- $A_{\text{burn}}^A$  regulates the timing that Alice and Bob can reveal  $pre_a$  and  $pre_b$ , respectively, and it is crucial to the proof of bounded-maximin fairness. We also add  $A_{\text{defuse}}^A$  to defuse the bomb of  $A_{\text{burn}}^A$ .
- $P_{\text{refund}}^A$  can be triggered by `ping` sent by Bob. This is to ensure that Bob can help Alice get refunded if the deposits are not finalized in time.

On BobChain, we modify the knowledge-coin exchange contract as follows:

- $C_{\text{burn}}^B$  can be triggered by  $pre_s$ ,  $pre_b$  and  $pre_c$ . The revelation of  $pre_c$  indicates that the deposits into the contracts on both blockchains are finalized in a timely manner. Thus, once the honest Alice posts  $pre_s$  to  $P_{\text{default}}^B$ , Bob should never try to get refunded by  $pre_b$ .
- $B_{\text{burn}}^B$  can be triggered by  $pre_c$  instead of  $pre_b$ . Because of the two-phase preparation, we need to ensure that Bob only posts  $pre_c$  after the deposits are finalized, which is crucial to the proof of bounded-maximin fairness.
- $P_{\text{refund}}^B$  can be triggered by `ping` sent by Alice. This step is crucial to for Alice to get refunded from BobChain if Bob drops out.

We give the complete specification of our contracts in Figure 3 and formally prove the full construction secure in Section 5.

## 3 Model

### 3.1 Blockchain and Smart Contracts

We model a *blockchain* as an append-only public ledger consisting of a sequence of ordered blocks. We consider an idealized mining process in which, at each block height  $t$ , a miner is selected with probability proportional to its mining power. This model captures both proof-of-work and proof-of-stake blockchains, where mining power corresponds to computational power and staked assets, respectively. The selected miner may then choose any subset of transactions from the pool of pending transactions and append a new block to the blockchain. Transactions in the block are executed in the order they appear.

We model a *smart contract* as an ideal functionality that can send and receive coins from players and the state of a smart contract is publicly observable. A smart contract may contain one or more *activation points*. Each transaction has a unique identifier and consists of: 1) an activation point of a smart contract, 2) a non-negative amount of coins, and 3) an arbitrary message. When a transaction is executed, the corresponding activation point of the smart contract is invoked, triggering the computation specified by the contract and possibly transferring coins. For simplicity, we use the phrase “A user sends a message `msg` to an activation point  $A$ ” to mean that the user sends a transaction to the activation point  $A$  with message `msg`. The balance of a smart contract is defined as the net amount of coins it has received minus those it has sent, and must remain non-negative.

We use the following style of pseudo-code to express smart contracts. We use `ping` to denote an empty message.

### A toy contract

- **Parameters:** time  $T$ . Alice deposits  $\$d_a$ , Bob deposits  $\$d_b$ .

$A_{\text{fast}}$ : On receiving `ping` from Alice: send  $\$d_b$  to Alice.

$A_{\text{wait}}$ : After  $T$ , on receiving `ping` from Alice: send  $\$d_a + \$d_b$  to Alice.

$B_{\text{other}}$ : On receiving `ping` from Bob: send  $\$d_a$  to Bob.

The capital letter in the contract defines the *type* of an activation point. All activation points of the same type are *mutually exclusive*. For example, if  $A_{\text{wait}}$  has been invoked, then neither  $A_{\text{fast}}$  nor  $A_{\text{wait}}$  can be invoked again. If an activation point is constrained to a specific time interval (e.g., after block height  $T$ ), then any attempt to invoke it outside that interval is considered invalid and is ignored. An activation point cannot be invoked if the contract’s balance is insufficient to cover the amount it is supposed to send. For example, if  $A_{\text{wait}}$  has been invoked,  $B_{\text{other}}$  cannot be invoked anymore.

The contract parameters specify the amounts of coins that must be deposited for the contract to become *active*. Once *all* required deposits are in place, the contract becomes active and its activation points can be invoked. Before the contract becomes active, each player is allowed to withdraw their own deposit if other players have not yet made the deposits. However, once the contract is active, the distribution of coins is only possible through the activation points.

## 3.2 Players and Strategy Spaces

There are three types of players in the model: Alice, Bob, and the miners. Alice and Bob are the two users who wish to exchange coins across different blockchains. We refer to them collectively as *users* to distinguish them from miners. All players are modeled as probabilistic polynomial-time (PPT) interactive Turing machines that can adaptively decide how to act based on the view in the protocol so far. We consider the following strategy spaces. Any player, including Alice, Bob, and miners, is allowed to do the following at any time:

1. **Post a transaction.** A player may post a transaction to the network at the beginning of any time step. We assume a synchronous network with zero delay. Transactions posted at time  $t$  are immediately visible to all players and miners. When miners decide which transactions to include in a block at time  $t$ , they can observe all transactions posted at that time.
2. **Create smart contracts.** A player may create an arbitrary smart contract and deposit any amount of coins into it. For example, a smart contract might specify: “If the blockchain state satisfies a particular predicate at some future time, transfer a specified amount of coins to a particular pseudonym,” where both the recipient and the amount may depend on the state of the blockchain.

In addition to the above, miners are further allowed the following:

3. **Block mining.** When a miner is selected to mine a block, it may include an arbitrary subset of the outstanding transactions into the block and order them arbitrarily. The miner may also generate and include new transactions of its own.

Alice or Bob can form a coalition with some miners. Within a coalition, all private information is shared among members. The strategy space of a coalition is defined as the union of the strategy spaces of its individual members.

Throughout the paper, except in Section 6, we do not explicitly model the process of coalition formation. Instead, we assume that all members within a coalition are bound together and act to

maximize the coalition’s joint utility. For simplicity, we ignore the transaction fee in our model. The bribery attack done by the transaction fee can be viewed as the transfer “within” the coalition, and does not affect the coalition’s utility. In Section 6, we analyze the metagame that models coalition formation.

### 3.3 Game-Theoretic Properties

**Cooperative Strategy Proofness.** The first property we consider is cooperative strategy proofness (CSP-fairness), which was formulated and adopted in prior works [PS17, CGL<sup>+</sup>18, WAS22, CS23, SCW23, CRS24, CMST22b]. Conceptually, a protocol satisfies CSP-fairness if following the protocol maximizes the expected utility of a coalition of strategic players, assuming that the remaining players follows the protocol honestly. Consequently, deviation is not profitable, and the honest protocol achieves a coalition-resistant Nash Equilibrium. The formal definition of CSP-fairness is as follows.

**Definition 3.1** (CSP fairness). A protocol satisfies  $\gamma$ -CSP-fairness iff the following holds. Let  $\mathcal{C}$  be any coalition that controls at most a  $\gamma \in [0, 1)$  fraction of the mining power, and possibly includes either Alice or Bob. Then, for any probabilistic polynomial-time (PPT) strategy  $S_{\mathcal{C}}$  of  $\mathcal{C}$ , there exists a negligible function  $\text{negl}(\cdot)$  such that except with  $\text{negl}(\lambda)$  probability, we have

$$\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, HS_{-\mathcal{C}}) \leq \text{util}^{\mathcal{C}}(HS_{\mathcal{C}}, HS_{-\mathcal{C}}), \quad (1)$$

where  $HS_{\mathcal{C}}$  denotes the honest strategy of  $\mathcal{C}$ ,  $HS_{-\mathcal{C}}$  denotes the honest strategy of anyone other than  $\mathcal{C}$ , and  $\text{util}^{\mathcal{C}}(X_{\mathcal{C}}, Y_{-\mathcal{C}})$  is the expected utility of the coalition  $\mathcal{C}$  when  $\mathcal{C}$  is executing strategy  $X$  and the remaining players (denoted by  $-\mathcal{C}$ ) execute strategy  $Y$ .<sup>1</sup>

**Dropout Resilience.** In atomic swap, we want to guarantee dropout resilience, which protects an honest player when the counterparty drops out from the protocol. In practice, a dropout can happen due to mistakes, misconfiguration, or unforeseen circumstances, e.g., Alice may lose her hardware wallet. Conceptually, a protocol satisfies dropout resilience if an honest player’s utility is always non-negative, even when the counterparty stops responding and drops out before the end of the protocol assuming at least  $1/\text{poly}(\lambda)$  fraction of the mining power is honest. The formal definition of dropout resilience is as follows.

**Definition 3.2** (Dropout resilience). A protocol is dropout resilient, iff as long as at least  $1/\text{poly}(\lambda)$  fraction of the mining power is honest, then with  $1 - \text{negl}(\lambda)$  probability, an honest Alice (or Bob) is guaranteed to have non-negative utility even when Bob (or Alice) is honest but drops out during the protocol’s execution.

We emphasize that our dropout resilience notion is very strong: we want it to hold even when many miners (up to  $1 - 1/\text{poly}(\lambda)$  fraction) are not necessarily playing honestly. In such a scenario, transactions may take polynomially long to confirm, and players may time out during the protocol and try to back out.

#### 3.3.1 New Definition: Bounded Maximin Fairness

While CSP-fairness assumes that players do not have incentives outside the present protocol, achieving resilience against such external incentives is very meaningful. In particular, such incentives may

---

<sup>1</sup>The formal definition of the utility function  $\text{util}$  is given in Section 3.4 in the context of atomic swap.

lead strategic players to adopt a broader class of strategies, including ones where players suffer loss within the protocol but are compensated by the external incentives. Given unbounded external incentives, strategic players can behave arbitrarily (like a “malicious” adversary using standard cryptographic terminology). The strongest game theoretic notion one can hope for is that honest players are not harmed even when others behave arbitrarily. This notion is called *maximin fairness* in some recent works [PS17, CGL<sup>+</sup>18, WAS22].

Maximin fairness, however, can be too stringent and challenging to satisfy. Thus, many real-world protocols that hope to incentivize honest behavior aim to provide incentive compatibility only when the parties’ external incentives are bounded. Specifically for atomic swap, it is unclear whether a protocol that achieves maximin fairness given user-miner coalition exists at all. Hence, we define a relaxed notion of fairness called *bounded maximin fairness*. It protects honest individuals from rational players who may have *arbitrary* but *bounded* external incentives. Although we focus on applying this notion to atomic swap, it would be interesting apply it more broadly to formally analyze real-world protocols that currently try to provide *heuristic* guarantees in the face of bounded external incentives.

**Defining bounded maximin fairness.** Imagine a set of players  $\mathcal{C}'$  who have external incentives that might entice them to deviate from honest behavior. We want to argue that even when their external incentives are *arbitrary*, then as long as the incentives are *bounded* and  $\mathcal{C}'$  is *rational*, any group of players without external incentives should feel safe to participate honestly — as long as they participate honestly, their utility will not be negative.

When defining *rationality* of the externally incentivized coalition, interesting technicalities arise due to the fact that we model players and contracts as PPT interactive Turing machines. It is tempting to assume that the externally incentivized coalition  $\mathcal{C}'$  uses the *optimal* strategy that maximizes its expected utility, and ensure no harm for any honest group or individual. However, this approach would make the possibly unrealistic assumption that  $\mathcal{C}'$  *knows* the optimal strategy based on the external incentive function. Indeed, finding such a strategy may be computationally hard, since the external incentive function can take an arbitrary form.

Instead, we will protect honest participants against any PPT strategy of the coalition  $\mathcal{C}'$  as long as the strategy is not *blatantly irrational*. A family of strategies  $\bar{\mathcal{R}}$  is called *blatantly irrational*, if for any strategy  $S \in \mathcal{R}$ , one can efficiently find another strategy  $S'$ , such that  $S'$  does strictly better than  $S$ .<sup>2</sup> We can often show that some strategies are *blatantly irrational* without finding the optimal strategy. Later in our formal proofs, we show that a class of strategies is *blatantly irrational*, since simple modifications of such strategies lead to better outcomes for  $\mathcal{C}'$ . We then show that as long as  $\mathcal{C}'$  does not adopt a *blatantly irrational* strategy, honest participants are protected.

We therefore devise the following definition which is parametrized by a strategy space  $\mathcal{R}$  that contains all possible PPT strategies except for the set of *blatantly irrational* strategies  $\bar{\mathcal{R}}$ .

**Definition 3.3** (Bounded maximin fairness). A protocol satisfies  $\alpha$ -bounded maximin fairness w.r.t. some strategy space  $\mathcal{R}$ , iff for any set of PPT players  $\mathcal{C}$  without external incentives, and any externally incentivized PPT coalition  $\mathcal{C}'$  that is disjoint from  $\mathcal{C}$ , controls at most  $\alpha$  fraction of the mining power, and uses any strategy  $S_{\mathcal{C}'} \in \mathcal{R}$ , there is a negligible function  $\text{negl}(\cdot)$  such that except with  $\text{negl}(\lambda)$  probability<sup>3</sup>, it holds that

$$\text{util}^{\mathcal{C}}(HS_{\mathcal{C}}, S_{\mathcal{C}'}, HS_{\mathcal{D}}) \geq 0,$$

<sup>2</sup>How we construct *blatantly irrational* strategies is given in Section 5.2.1, and how one can find a better strategy efficiently is explained in Lemma 5.7 and Lemma 5.10.

<sup>3</sup>The negligible failure probability in our proofs arises due to the following bad events: 1) either the hash is inverted, or 2) honest miners have never mined a block even after polynomially many blocks have been mined.

where  $\mathcal{D}$  denotes all players not in  $\mathcal{C} \cup \mathcal{C}'$ ,  $HS_{\mathcal{C}}$  and  $HS_{\mathcal{D}}$  denote the honest strategy of  $\mathcal{C}$  and  $\mathcal{D}$ , respectively, and  $\text{util}^{\mathcal{C}}(HS_{\mathcal{C}}, S_{\mathcal{C}'}, HS_{\mathcal{D}})$  denotes the utility of  $\mathcal{C}$  given that the strategy  $(HS_{\mathcal{C}}, S_{\mathcal{C}'}, HS_{\mathcal{D}})$  is adopted.

**Regarding external incentives.** As mentioned above, our modeling of external incentives is general and captures external incentives of any form. Precisely, any side contract where money is redistributed to players of the present protocol is considered as external incentive if **(1)** the contract is *pre-existing*, i.e., created before the start of the present protocol; or **(2)** the contract is created at any time, and an *outsider*, i.e., not a player of the present protocol, deposited money into it. As an example of the former, imagine that Alice was involved in some bet prior to the atomic swap protocol, and the bet depends on the state of a future block. Because the outcome of the atomic swap protocol may affect the blockchain state, Alice may be incentivized to manipulate the state even if doing so lead to a loss *within* the atomic swap protocol. As an example of the latter, imagine that Alice is Mallory's competitor, and Mallory is offering external incentives for anyone who can cause financial loss to Alice. Without loss of generality, we consider non-negative external incentives. Further, for pre-existing smart contracts, we do not distinguish between those funded by the participants in our protocol and those funded by external parties — gains from all pre-existing smart contracts are considered external.

### 3.4 Atomic Swap

An atomic swap allows two mutually distrustful parties to exchange coins across different blockchains without relying on a trusted third party. Consider two parties, Alice and Bob. Bob owns  $x_b$  coins on **BobChain** (denoted  $\mathbb{B}x_b$ ), and Alice owns  $x_a$  coins on **AliceChain** (denoted  $\mathbb{A}x_a$ ). The goal is for Bob to exchange his  $\mathbb{B}x_b$  for Alice's  $\mathbb{A}x_a$ .

In this work, we consider three kinds of strategic players: 1) Alice-miner coalition (or Alice alone); 2) Bob-miner coalition (or Bob alone); and 3) miner-only coalition. Let  $\$AV(\cdot)$  denote the valuation function of Alice (or the Alice-miner coalition), defined as:

$$\$AV(\mathbb{B}x_b + \mathbb{A}x_a) = \$v_a^B \cdot x_b + \$v_a^A \cdot x_a,$$

where  $\$v_a^B \geq 0$  and  $\$v_a^A \geq 0$  represent the value Alice places on each coin on **BobChain** and **AliceChain**, respectively. Similarly, let  $\$BV(\cdot)$  denote the valuation function of Bob (or the Bob-miner coalition), and  $\$MV(\cdot)$  denote that of the miner-only coalition. We make the assumption:  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a) > 0$ ,  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b) > 0$ , which justifies why Alice wants to exchange her  $\mathbb{A}x_a$  with Bob's  $\mathbb{B}x_b$ , and vice versa.

Let  $\mathcal{C}$  be any subset of players, with  $S_{\mathcal{C}}$  and  $S'_{-\mathcal{C}}$  denoting the strategies of  $\mathcal{C}$  and  $-\mathcal{C}$ , respectively. Let  $\mathbb{A}d_a^A, \mathbb{B}d_a^B \geq 0$  be the amounts deposited by Alice (or the Alice-miner coalition) into the respective smart contracts, and  $\mathbb{A}r_a^A, \mathbb{B}r_a^B \geq 0$  be the amounts received during protocol execution. Let  $\$e_a(\dots) \geq 0$  represent the external incentives for Alice or the Alice-miner coalition, where the value may depend arbitrarily on the blockchain state. We define the utility  $\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}})$  when  $\mathcal{C}$  is Alice or Alice-miner coalition as follows:

$$\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}}) = \$AV(\mathbb{A}r_a^A - \mathbb{A}d_a^A + \mathbb{B}r_a^B - \mathbb{B}d_a^B) + \$e_a(\dots).$$

We define  $\mathbb{A}d_b^A, \mathbb{B}d_b^B, \mathbb{A}r_b^A, \mathbb{B}r_b^B, \$e_b(\dots) \geq 0$  analogously for Bob (or the Bob-miner coalition), and  $\mathbb{A}r_m^A, \mathbb{B}r_m^B, \mathbb{A}d_m^A, \mathbb{B}d_m^B, \$e_m(\dots) \geq 0$  for the miner-only coalition. The utility  $\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}})$  when  $\mathcal{C}$  consists of Bob or a Bob-miner coalition is defined as

$$\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}}) = \$BV(\mathbb{A}r_b^A - \mathbb{A}d_b^A + \mathbb{B}r_b^B - \mathbb{B}d_b^B) + \$e_b(\dots),$$

and the utility  $\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}})$  when  $\mathcal{C}$  is a miner-only coalition is defined as

$$\text{util}^{\mathcal{C}}(S_{\mathcal{C}}, S'_{-\mathcal{C}}) = \$\text{MV}(\mathbb{A}r_m^{\mathbb{A}} - \mathbb{A}d_m^{\mathbb{A}} + \mathbb{B}r_m^{\mathbb{B}} - \mathbb{B}d_m^{\mathbb{B}}) + \$e_m(\dots).$$

## 4 Secure Atomic Swap

We now present the smart contracts and the protocol of our atomic swap. Say Alice owns  $x_a$  coins on **AliceChain** (denoted  $\mathbb{A}x_a$ ), and Bob owns  $x_b$  coins on **BobChain** (denoted  $\mathbb{B}x_b$ ). Bob wants to exchange his  $\mathbb{B}x_b$  for Alice's  $\mathbb{A}x_a$ . Either Alice or Bob can be a strategic player, and collude with the miners on both chains.  $\$E$  is the upper bound of the external incentive of the coalition, and  $\alpha \in [0, 1]$  is the fraction of the mining power controlled by the coalition.

To initiate an atomic swap, Alice samples  $pre_s \leftarrow \{0, 1\}^{\lambda}$ ,  $pre_a \leftarrow \{0, 1\}^{\lambda}$ , and Bob samples  $pre_b \leftarrow \{0, 1\}^{\lambda}$  and  $pre_c \leftarrow \{0, 1\}^{\lambda}$ , where  $\lambda$  is the security parameter. Alice and Bob choose the parameters that respect the constraints in Figure 2. Then, Alice and Bob deploy two contracts:  $\text{CONTRACT}^{\mathbb{A}}$  and  $\text{CONTRACT}^{\mathbb{B}}$  specified in Figure 3. We rename the block height where  $\text{CONTRACT}^{\mathbb{A}}$  is deployed on **AliceChain** to be **AliceChain** time 0 (only deployed, Alice and Bob have not deposited their coins yet). Similarly, **BobChain** time 0 is the block height where  $\text{CONTRACT}^{\mathbb{B}}$  is deployed on **BobChain**.

We give the protocols specifying the honest behaviors of Alice, Bob, and miners in Figure 4, Figure 5, and Figure 6, respectively. This atomic swap protocol satisfies CSP-fairness, bounded maximin fairness, and dropout resilience:

**Theorem 4.1.** *Suppose that the hash function  $H(\cdot)$  is a one-way function, and the choice of the parameters satisfy the constraints in Figure 2. Then, the following statements hold:*

- For any  $\gamma \in [0, 1 - 1/\text{poly}(\lambda)]$ , if the parameters further satisfy  $\gamma^{\tau^{\mathbb{A}}} \leq \frac{\mathbb{A}c_a^{\mathbb{A}}}{\mathbb{A}c_a^{\mathbb{A}} + \mathbb{A}x_a}$  and  $\gamma^{\tau^{\mathbb{B}}} \leq \frac{\mathbb{B}c_b^{\mathbb{B}}}{\mathbb{B}c_b^{\mathbb{B}} + \mathbb{B}x_b}$ , then atomic swap protocol satisfies  $\gamma$ -CSP-fairness.
- If  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ , then atomic swap protocol satisfies  $\alpha$ -bounded maximin fairness against external incentives.
- If all players are PPT, then atomic swap protocol is dropout resilient.

The formal proof of Theorem 4.1 is given in Section 5. In Section 6, we analyze the metagame that captures the formation of the coalition. We show that our atomic swap protocol disincentivizes 100% of the miners to collude with the strategic players, and thus justifies the assumption of having  $1/\text{poly}(\lambda)$  fraction of honest miners.

In addition to the game-theoretic properties above, we also show that honest players can get their collateral back in a timely manner as specified in the following theorem.

**Theorem 4.2.** *Suppose Alice, Bob, and all miners are honest. Alice gets  $\mathbb{B}x_b$ , Bob gets  $\mathbb{A}x_a$ , and both parties get their collateral back in 3 **BobChain** time plus 2 **AliceChain** time.*

*Proof.* If Alice and Bob are both honest, they will send the deposit transactions to  $\text{CONTRACT}^{\mathbb{B}}$  and  $\text{CONTRACT}^{\mathbb{A}}$  at **BobChain** time  $t = 0$ . Because all the miners are honest,  $\text{CONTRACT}^{\mathbb{B}}$  and  $\text{CONTRACT}^{\mathbb{A}}$  both enter the execution phase no later than one block on **BobChain** or one block on **AliceChain** is mined. As soon as  $\text{CONTRACT}^{\mathbb{B}}$  and  $\text{CONTRACT}^{\mathbb{A}}$  both enter the execution phase, Bob sends ping to  $B_{\text{defuse}}^{\mathbb{B}}$ . Because all the miners are honest,  $B_{\text{defuse}}^{\mathbb{B}}$  will be activated no later than one block on **BobChain** is mined. As soon as  $B_{\text{defuse}}^{\mathbb{B}}$  is activated, Bob sends  $pre_c$  to  $P_{\text{default}}^{\mathbb{B}}$ .

**Constraints for Contract<sup>B</sup> (on BobChain):**

- $h_s = H(pre_s)$ ,  $h_b = H(pre_b)$  and  $h_c = H(pre_c)$ .
- $T_1^B > T_0^B > T^B > 0$ .
- $\mathbb{B}\epsilon^B > \mathbb{B}0$ ,  $\mathbb{B}c_a^B > \mathbb{B}\epsilon^B$ , and  $\mathbb{B}c_b^B > \mathbb{B}\epsilon^B$ .
- $\$AV(\mathbb{B}c_a^B) > \frac{\$AV(\mathbb{A}x_a + \alpha \mathbb{B}x_b) + \$E}{1-\alpha}$  and  
 $\$BV(\mathbb{B}c_b^B) > \frac{\$BV(\mathbb{A}x_a + \alpha \mathbb{B}x_b) + \$E}{1-\alpha}$

**Constraints for Contract<sup>A</sup> (on AliceChain):**

- $h_s = H(pre_s)$  and  $h_a = H(pre_a)$ .
- $T_1^A > T_0^A > 0$ .
- AliceChain time  $T_1^A >$  BobChain time  $T_1^B$ , i.e., the AliceChain block of length  $T_1^A$  is mined after the BobChain block of length  $T_1^B$ .<sup>a</sup>
- $\mathbb{A}\epsilon^A > \mathbb{A}0$ ,  $\mathbb{A}c_a^A > \mathbb{A}\epsilon^A$ , and  $\mathbb{A}c_b^A > \mathbb{A}\epsilon^A$ .
- $\$AV(\mathbb{A}c_a^A) > \frac{\$AV(\mathbb{B}x_b + \alpha \mathbb{A}x_a) + \$E}{1-\alpha}$  and  
 $\$BV(\mathbb{A}c_b^A) > \frac{\$BV(\mathbb{B}x_b + \alpha \mathbb{A}x_a) + \$E}{1-\alpha}$ .

**Choice of timeouts:**

- $\tau^B \geq 1$ ,  $\tau^A \geq 1$ .

---

<sup>a</sup>In practice, this constraint should be respected except with negligible probability despite the variance in inter-block times.

Figure 2: Parameter constraints for atomic swap.  $H(\cdot)$  is a cryptographic hash function. All times are expressed in the time of the respective chain.

Because we assume the network delay is zero, Alice and Bob both enter the execution phase when Bob sends  $pre_c$  to  $P_{\text{default}}^B$ .

When Alice enters the execution phase, she sends  $pre_s$  to  $P_{\text{default}}^B$  immediately. Because all the miners are honest,  $P_{\text{default}}^B$  will be activated no later than one block on BobChain is mined. As soon as  $P_{\text{default}}^B$  is activated, Alice sends ping to  $P_{\text{default}}^A$ . Again, because all the miners are honest,  $P_{\text{default}}^A$  will be activated no later than one block on AliceChain is mined. When  $P_{\text{default}}^B$  and  $P_{\text{default}}^A$  are activated, Alice gets  $\mathbb{B}x_b$ , Bob gets  $\mathbb{A}x_a$ , and both parties get their collateral back.  $\square$

### Contract<sup>B</sup> (on BobChain)

/\* parametrized with  $(h_s, h_b, h_c, T_1^B, \tau^B, \mathbb{B}x_b, \mathbb{B}c_b^B, \mathbb{B}\epsilon^B)^*$  \*/

**Deposits:** Bob deposits  $\mathbb{B}x_b + \mathbb{B}c_b^B$ , and Alice deposits  $\mathbb{B}c_a^B$ . Once both parties have deposited the required amount, the contract becomes active.

$B_{\text{defuse}}^B$ : On receiving ping from anyone, do nothing.  $\triangleright B_{\text{defuse}}^B$  is to invalidate  $B_{\text{burn}}^B$ .

$B_{\text{burn}}^B$ : On receiving  $z$  from anyone  $P$  such that  $H(z) = h_c$ , send  $\mathbb{B}\epsilon^B$  to player  $P$ . All remaining coins are burnt.

$P_{\text{default}}^B$ : On receive  $z_1$  from Alice such that  $H(z_1) = h_s$  and  $z_2$  from Bob such that  $H(z_2) = h_c$ , send  $\mathbb{B}x_b + \mathbb{B}c_a^B$  to Alice and  $\mathbb{B}c_b^B$  to Bob.

$P_{\text{refund}}^B$ : Time  $T_1^B$  or greater: On receive  $z$  from Bob such that  $H(z) = h_b$  or on receiving ping from Alice, do nothing.

$C_{\text{refund}}^B$ : At least  $\tau^B$  after  $P_{\text{refund}}^B$  is activated: on receiving ping from anyone, send  $\mathbb{B}x_b + \mathbb{B}c_b^B$  to Bob and  $\mathbb{B}c_a^B$  to Alice.

$C_{\text{burn}}^B$ : On receive  $(z_1, z_2, z_3)$  from anyone  $P$  such that  $H(z_1) = h_s$ ,  $H(z_2) = h_b$ , and  $H(z_3) = h_c$  send  $\mathbb{B}\epsilon^B$  to player  $P$ . All remaining coins are burnt.

### Contract<sup>A</sup> (on AliceChain)

/\* parametrized with  $(h_s, h_a, T_1^A, \tau^A, \mathbb{A}x_a, \mathbb{A}c_a^A, \mathbb{A}\epsilon^A)^*$  \*/

**Deposits:** Alice deposits  $\mathbb{A}x_a + \mathbb{A}c_a^A$ , and Bob deposits  $\mathbb{A}c_b^A$ . Once both parties have deposited the required amount, the contract becomes active.

$A_{\text{defuse}}^A$ : Time  $T_1^A$  or greater: on receiving ping from Alice or Bob, do nothing.  $\triangleright A_{\text{defuse}}^A$  is to invalidate  $A_{\text{burn}}^A$ .

$A_{\text{burn}}^A$ : On receiving  $z_1$  from anyone  $P$  such that  $H(z_1) = h_a$ , or on receiving  $z_2$  from anyone  $P$  such that  $H(z_2) = h_b$ , send  $\mathbb{A}\epsilon^A$  to player  $P$ . All remaining coins are burnt.

$P_{\text{default}}^A$ : On receiving  $z$  from Bob such that  $H(z) = h_s$  or on receiving ping from Alice, send  $\mathbb{A}x_a + \mathbb{A}c_b^A$  to Bob and send  $\mathbb{A}c_a^A$  to Alice.

$P_{\text{refund}}^A$ : Time  $T_1^A$  or greater: on receiving  $z$  from Alice such that  $H(z) = h_a$  or on receiving ping from Bob, do nothing.

$C_{\text{refund}}^A$ : At least  $\tau^A$  after  $P_{\text{refund}}^A$  is activated: on receiving ping from anyone, send  $\mathbb{A}x_a + \mathbb{A}c_a^A$  to Alice and  $\mathbb{A}c_b^A$  to Bob.

$C_{\text{burn}}^A$ : On receiving  $(z_1, z_2)$  from anyone  $P$  such that  $(z_1) = h_s$  and  $H(z_2) = h_a$ , or on receiving  $(z_2, z_3)$  from anyone  $P$  such that  $H(z_2) = h_a$  and  $H(z_3) = h_b$ , send  $\mathbb{A}\epsilon^A$  to player  $P$ . All remaining coins are burnt.

Figure 3: Smart contracts for atomic swap. Activation points of the same type are mutually exclusive. Activation points can be triggered only after both parties have deposited.

## Atomic Swap Protocol — Alice

### Preparation Phase:

1. When both  $\text{CONTRACT}^A$  and  $\text{CONTRACT}^B$  have been deployed on the respective chains, Alice sends the deposit  $\mathbb{A}x_a + \mathbb{A}c_a^A$  to  $\text{CONTRACT}^A$ ; and the collateral  $\mathbb{B}c_a^B$  to  $\text{CONTRACT}^B$ .
2. Wait until one of the following happens:
  - Either  $\text{CONTRACT}^B$  or  $\text{CONTRACT}^A$  has not been active, and it is at least BobChain time  $T^B$ : Alice enters the abort phase.
  - Bob has not sent  $pre_c$  to  $P_{\text{default}}^B$ , and it is at least BobChain time  $T_0^B$ : Alice enters the abort phase.
  - Bob sent  $pre_c$  to  $P_{\text{default}}^B$  and it is before BobChain time  $T_0^B$ : Alice enters the execution phase.

### Execution phase:

1. Alice sends  $pre_s$  to  $P_{\text{default}}^B$ . As soon as  $P_{\text{default}}^B$  has been activated, Alice sends ping to  $P_{\text{default}}^A$ .
2. If  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Alice sends ping to  $C_{\text{refund}}^B$ . (Note that as soon as  $C_{\text{refund}}^B$  is activated, Bob sends ping to  $P_{\text{refund}}^A$ .)
3. If  $\tau^A$  AliceChain time has passed since activating  $P_{\text{refund}}^A$ , Alice sends ping to  $C_{\text{refund}}^A$ .

### Abort Phase:

1. If  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.) has not been active, Alice withdraws her deposit from  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.).
2. At AliceChain time  $T_0^A$ , Alice sends ping to  $P_{\text{refund}}^B$  and ping to  $A_{\text{defuse}}^A$ .
3. If Bob has not sent ping to  $P_{\text{refund}}^A$  by AliceChain time  $T_1^A$ , Alice waits until  $A_{\text{defuse}}^A$  is activated and sends  $pre_a$  to  $P_{\text{refund}}^A$ .
4. If  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated, Alice sends ping to  $C_{\text{refund}}^A$ ; similarly, if  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Alice sends ping to  $C_{\text{refund}}^B$ .

### Ignore all other events.

Figure 4: Atomic swap protocol for Alice.

## Atomic Swap Protocol — Bob

### Preparation Phase:

1. When both  $\text{CONTRACT}^A$  and  $\text{CONTRACT}^B$  have been deployed on the respective chains, Bob sends the deposit transaction of  $\mathbb{B}x_b + \mathbb{B}c_b^B$  to  $\text{CONTRACT}^B$  and sends the collateral transaction of  $\mathbb{A}c_b^A$  to  $\text{CONTRACT}^A$ .
2. Wait until one of the following happens:
  - Both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  are active: Bob sends ping to  $B_{\text{defuse}}^B$ . As soon as  $B_{\text{defuse}}^B$  is activated, Bob sends  $pre_c$  to  $P_{\text{default}}^B$  and enters the execution phase.
  - Either  $\text{CONTRACT}^B$  or  $\text{CONTRACT}^A$  has not been active, and it is at least BobChain time  $T^B$ : Bob enters the abort phase;

### Execution phase:

1. Wait until one of the following happens:
  - Alice sent  $pre_s$  to  $P_{\text{default}}^B$ : Bob sends  $pre_s$  to  $P_{\text{default}}^A$ .
  - Alice has not sent  $pre_s$  to  $P_{\text{default}}^B$ , and it is at least BobChain time  $T_1^B$ : Bob sends ping to  $A_{\text{defuse}}^A$ . As soon as  $A_{\text{defuse}}^A$  is activated, Bob sends  $pre_b$  to  $P_{\text{refund}}^B$ .
2. If  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Bob sends ping to  $C_{\text{refund}}^B$ . As soon as  $C_{\text{refund}}^B$  is activated, Bob sends ping to  $P_{\text{refund}}^A$ .
3. If  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated, Bob sends ping to  $C_{\text{refund}}^A$ .

### Abort Phase:

1. If  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.) has not been active, Bob withdraws his deposit from  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.).
2. At AliceChain time  $T_0^A$ , Bob sends ping to  $P_{\text{refund}}^A$  and ping to  $A_{\text{defuse}}^A$ .
3. If Alice has not sent ping to  $P_{\text{refund}}^B$  by AliceChain time  $T_1^A$ , Bob waits until  $A_{\text{defuse}}^A$  is activated and sends  $pre_b$  to  $P_{\text{refund}}^B$ .
4. If  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated, Bob sends ping to  $C_{\text{refund}}^A$ ; similarly, if  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Bob sends ping to  $C_{\text{refund}}^B$ .

**Ignore all other events.**

Figure 5: Atomic swap protocol for Bob.

### Atomic Swap Protocol — Miner

- The miner watches all transactions posted to  $B_{\text{defuse}}^B$ ,  $B_{\text{burn}}^B$ ,  $P_{\text{default}}^B$ ,  $P_{\text{refund}}^B$ ,  $C_{\text{refund}}^B$ ,  $C_{\text{burn}}^B$ ,  $A_{\text{defuse}}^A$ ,  $A_{\text{burn}}^A$ ,  $P_{\text{default}}^A$ ,  $P_{\text{refund}}^A$ ,  $C_{\text{refund}}^A$ , and  $C_{\text{burn}}^A$  (i.e., all the activation points for both contracts), to see if they contain a valid  $\text{pre}_s$ ,  $\text{pre}_a$ ,  $\text{pre}_b$ , and  $\text{pre}_c$ .
- If  $C_{\text{refund}}^B$  has not been activated, as soon as the miner has observed  $\text{pre}_s$ ,  $\text{pre}_b$  and  $\text{pre}_c$ , it posts  $(\text{pre}_s, \text{pre}_b, \text{pre}_c)$  to  $C_{\text{burn}}^B$ . Similarly, if  $C_{\text{refund}}^A$  has not been activated, as soon as the miner has observed both  $\text{pre}_a$  and  $\text{pre}_s$ , it posts  $(\text{pre}_a, \text{pre}_s)$  to  $C_{\text{burn}}^A$ ; as soon as the miner has observed  $\text{pre}_a$  and  $\text{pre}_b$ , it posts  $(\text{pre}_a, \text{pre}_b)$  to  $C_{\text{burn}}^A$ .
- If  $B_{\text{defuse}}^B$  has not been activated, as soon as the miner observes  $\text{pre}_c$ , it posts  $\text{pre}_c$  to  $B_{\text{burn}}^B$ . If  $A_{\text{defuse}}^A$  has not been activated, as soon as miner observes  $\text{pre}_a$  or  $\text{pre}_b$ , it posts  $\text{pre}_a$  or  $\text{pre}_b$  to  $A_{\text{burn}}^A$ .
- Whenever the miner mines a block, it always includes its own transactions ahead of others.

Figure 6: Atomic swap protocol for miners.

## 5 Proof of Theorem 4.1

### 5.1 Achieving CSP Fairness

This subsection is dedicated to proving CSP-fairness which formally stated as follows.

**Theorem 5.1** (CSP fairness). *Suppose that the hash function  $H(\cdot)$  is a one-way function. Suppose the choice of the parameters satisfy the constraints in Figure 2, and further satisfy  $\gamma^{\tau^A} \leq \frac{\mathbb{A}c_a^A}{\mathbb{A}c_a^A + \mathbb{A}x_a}$  and  $\gamma^{\tau^B} \leq \frac{\mathbb{B}c_b^B}{\mathbb{B}c_b^B + \mathbb{B}x_b}$ . Then, the atomic swap protocol satisfies  $\gamma$ -CSP-fairness.*

Before proving Theorem 5.1, we define the following events.

- **Normal<sup>B</sup>**:  $P_{\text{default}}^B$  is activated.
- **Refund<sup>B</sup>**: either  $(P_{\text{refund}}^B + C_{\text{refund}}^B)$  are activated, or one of Alice and Bob withdraws their deposits from  $\text{CONTRACT}^B$  successfully before  $\text{CONTRACT}^B$  becomes active.
- **Burn<sup>B</sup>**:  $B_{\text{burn}}^B$  or  $C_{\text{burn}}^B$  is activated.
- **Normal<sup>A</sup>**:  $P_{\text{default}}^A$  is activated.
- **Refund<sup>A</sup>**: either  $(P_{\text{refund}}^A + C_{\text{refund}}^A)$  are activated, or one of Alice and Bob withdraws their deposits from  $\text{CONTRACT}^A$  successfully before  $\text{CONTRACT}^A$  becomes active.
- **Burn<sup>A</sup>**:  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated.

Normally, when Alice and Bob follow the protocol,  $P_{\text{default}}^B$  and  $P_{\text{default}}^A$  will be activated, and they exchange the coins successfully. However, if one of the parties drops out, the other party will trigger  $(P_{\text{refund}}^B + C_{\text{refund}}^B)$  and  $(P_{\text{refund}}^A + C_{\text{refund}}^A)$  to get refunded. Finally,  $B_{\text{burn}}^B$ ,  $C_{\text{burn}}^B$ ,  $A_{\text{burn}}^A$  and  $C_{\text{burn}}^A$  are the bombs, and both Alice and Bob lose their collateral when a bomb is triggered.

**Lemma 5.2.** *Suppose the parameters are set according to Figure 2. Then, the following statements hold.*

- Suppose the coalition  $\mathcal{A}$  consists of Alice and an arbitrary  $\gamma \in [0, 1]$  fraction of the mining power. The utility of  $\mathcal{A}$  can be more than the honest case, that is,  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a)$ , only if **Normal<sup>B</sup>** and **Refund<sup>A</sup>** both happen.
- Suppose the coalition  $\mathcal{B}$  consists of Bob and an arbitrary  $\gamma \in [0, 1]$  fraction of the mining power. The utility of  $\mathcal{B}$  can be more than the honest case, that is,  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b)$ , only if **Refund<sup>B</sup>** and **Normal<sup>A</sup>** both happen.

*Proof.* Notice that if any of **Normal<sup>B</sup>**, **Refund<sup>B</sup>** and **Burn<sup>B</sup>** happens, no coin is left in BobChain, so no one can get more coin from BobChain anymore. Thus, consider all possible cases, including none of **Normal<sup>B</sup>**, **Refund<sup>B</sup>** and **Burn<sup>B</sup>** happens, we have the following table.

which is activated	net profit of Alice's coalition	net profit of Bob's coalition
none	$-\mathbb{B}c_a^B$	$-\mathbb{B}x_b - \mathbb{B}c_b^B$
Normal <sup>B</sup>	$\mathbb{B}x_b$	$-\mathbb{B}x_b$
Refund <sup>B</sup>	0	0
Burn <sup>B</sup>	$\leq \mathbb{B}\epsilon^B - \mathbb{B}c_a^B$	$\leq \mathbb{B}\epsilon^B - \mathbb{B}x_b - \mathbb{B}c_b^B$

Table 1: The net profit of Bob's coalition from BobChain.

Similarly, if any of  $\text{Normal}^A$ ,  $\text{Refund}^A$  and  $\text{Burn}^A$  happens, no coin is left in AliceChain, so no one can get more coin from AliceChain anymore. Thus, consider all possible cases, including none of  $\text{Normal}^A$ ,  $\text{Refund}^A$  and  $\text{Burn}^A$  happens, we have the following table.

which is activated	net profit of Alice's coalition	net profit of Bob's coalition
none	$-\mathbb{A}x_a - \mathbb{A}c_a^A$	$-\mathbb{A}c_b^A$
$\text{Normal}^A$	$-\mathbb{A}x_a$	$\mathbb{A}x_a$
$\text{Refund}^A$	0	0
$\text{Burn}^A$	$\leq \mathbb{A}\epsilon^A - \mathbb{A}x_a - \mathbb{A}c_a^A$	$\leq \mathbb{A}\epsilon^A - \mathbb{A}c_b^A$

Table 2: The net profit of Alice's coalition and Bob's coalition from AliceChain.

**Alice-miner coalition.** Suppose the coalition  $\mathcal{A}$  consists of Alice and an arbitrary  $\gamma \in [0, 1]$  fraction of the mining power. If  $\mathcal{A}$  follows the protocol,  $\text{Normal}^B$  and  $\text{Normal}^A$  will happen, and the utility of  $\mathcal{C}$  is  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a) > 0$ . When  $\text{Normal}^B$  and  $\text{Refund}^A$  both happen,  $\mathcal{A}$ 's utility is  $\$AV(\mathbb{B}x_b)$ . Now, we will show that this is the only possible scenario for  $\mathcal{A}$ 's utility to exceed the honest case. In other words, if either  $\text{Normal}^B$  or  $\text{Refund}^A$  does not happen, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a)$ . There are two cases.

- **Case 1:  $\text{Normal}^B$  does not happen.** Because  $\mathbb{B}c_a^B > \mathbb{B}\epsilon^B$ , the net profit from BobChain is at most 0 if  $\text{Normal}^B$  does not happen. Then, because  $\mathbb{A}x_a > \mathbb{A}\epsilon^A$ , we have  $\mathbb{A}\epsilon^A - \mathbb{A}x_a - \mathbb{A}c_a^A < 0$ . Thus, the net profit from AliceChain is also at most 0. Consequently, the utility of  $\mathcal{C}$  is at most zero, which is less than  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a)$ .
- **Case 2:  $\text{Refund}^A$  does not happen.** Because  $\mathbb{A}c_a^A > \mathbb{A}\epsilon^A$ , we have  $\mathbb{A}\epsilon^A - \mathbb{A}x_a - \mathbb{A}c_a^A < -\mathbb{A}x_a$ . Thus, assuming  $\text{Refund}^A$  does not happen, the net profit from AliceChain is at most  $-\mathbb{A}x_a$ . However, the net profit from BobChain is at most  $\mathbb{B}x_b$ . Thus, the utility of  $\mathcal{C}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a)$ , which is the same as the honest case.

**Bob-miner coalition.** Using a completely symmetric proof, we can show that the only way for a Bob-miner coalition's utility to exceed the honest case is when  $\text{Refund}^B$  and  $\text{Normal}^A$  both happen.  $\square$

**Lemma 5.3** (Alice-miner coalition). *Suppose that the hash function  $H(\cdot)$  is a one-way function. Let  $\mathcal{A}$  be any coalition that consists of Alice and  $\gamma$  fraction of mining power. Then, as long as  $\gamma^A \leq \frac{\mathbb{A}c_a^A}{\mathbb{A}c_a^A + \mathbb{A}x_a}$ , for any PPT coalition strategy  $S_{\mathcal{A}}$ , except with negligible probability, it must be*

$$\text{util}^A(S_{\mathcal{A}}, HS_{-\mathcal{A}}) \leq \text{util}^A(HS_{\mathcal{A}}, HS_{-\mathcal{A}}),$$

where  $HS_{-\mathcal{A}}$  denotes the honest strategy for everyone not in  $\mathcal{A}$ .

*Proof.* Recall that the utility of  $\mathcal{A}$  is  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a) > 0$  under an honest execution. Now, suppose  $\mathcal{A}$  may deviate from the protocol. We may assume that the coalition does not post any new smart contract on the fly and deposit money into it — if it did so, it cannot recover more than its deposit since any player not in  $\mathcal{A}$  will not invoke the smart contract. We analyze the possible cases depending on which phase Bob enters.

**Bob enters the abort phase.** Because  $\mathbb{A}x_a > \mathbb{A}\epsilon^A$ , the net profit of  $\mathcal{A}$  from AliceChain is at most zero, no matter whether  $\text{CONTRACT}^A$  is active or not. We will show that the net profit of

$\mathcal{A}$  from BobChain is also at most zero except with negligible probability. If  $\text{CONTRACT}^B$  never becomes active, the net profit of  $\mathcal{A}$  from BobChain is at most zero. Now, assume  $\text{CONTRACT}^B$  becomes active. When Bob enters the abort phase, he never sends any transaction containing  $pre_c$ . Ignoring the negligible probability that  $\mathcal{A}$  finds  $pre_c$  by itself,  $B_{\text{burn}}^B$ ,  $P_{\text{default}}^B$ , and  $C_{\text{burn}}^B$  can never be activated. Because Alice does not get any coin from  $B_{\text{defuse}}^B$ ,  $P_{\text{refund}}^B$  or  $C_{\text{refund}}^B$ , the net profit of  $\mathcal{A}$  from BobChain is at most zero.

To sum up, except with negligible probability, the utility of  $\mathcal{A}$  is at most zero, which is less than the honest case.

**Bob enters the execution phase.** If Bob enters the execution phase, both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  must be active. By Lemma 5.2, the utility of  $\mathcal{A}$  can exceed the honest case only when  $\text{Normal}^B$  and  $\text{Refund}^A$  both happen, so we assume it is the case. Because  $\text{CONTRACT}^A$  is active, for  $\text{Refund}^A$  to happen,  $P_{\text{refund}}^A$  must be activated. When Bob enters the execution phase,  $P_{\text{refund}}^A$  can be activated only either 1) by Bob sending ping to  $P_{\text{refund}}^A$  after  $C_{\text{refund}}^B$  has been activated, or 2) by Alice sending  $pre_a$  to  $P_{\text{refund}}^A$ . Consider the first scenario. Since  $C_{\text{refund}}^B$  has been activated, Alice cannot get any money from  $\text{CONTRACT}^B$ . However, from  $\text{CONTRACT}^A$ , Alice can get at most zero. Thus, the utility of  $\mathcal{A}$  is less than the honest case.

Now consider the second case. Suppose that  $P_{\text{refund}}^A$  is activated at AliceChain time  $t^* \geq T_1^A$ , so  $pre_a$  is publicly known after AliceChain time  $t^*$ . By assumption,  $\text{Normal}^B$  happens, so  $P_{\text{default}}^B$  must be activated. In this case,  $\mathcal{A}$  has to send  $pre_s$  to  $P_{\text{default}}^B$ .

- *Case 1:  $\mathcal{A}$  sends  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ .* Since BobChain time  $T_1^B$  is earlier than AliceChain time  $T_1^A$ ,  $pre_s$  and  $pre_a$  are both publicly known at AliceChain time  $t^*$ . Thus, during AliceChain time  $(t^*, t^* + \tau^A]$ , any honest miner will activate  $C_{\text{burn}}^A$  if it wins a block. We say  $\mathcal{A}$  loses the race if a non-colluding miner mines a new block during AliceChain time  $(t^*, t^* + \tau^A]$ . Otherwise, we say  $\mathcal{A}$  wins the race. If  $\mathcal{A}$  loses the race, it gets nothing from  $C_{\text{refund}}^A$  or  $C_{\text{burn}}^A$ , and its utility is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A)$ . Else if  $\mathcal{A}$  wins the race, then its utility is at most  $\$AV(\mathbb{B}x_b)$ , which can be achieved by activating  $P_{\text{refund}}^A$ ,  $C_{\text{refund}}^A$  and  $P_{\text{default}}^B$ . The probability  $p$  that  $\mathcal{A}$  wins the race is upper bounded by  $p \leq \gamma^{\tau^A}$ . Therefore, the expected utility of  $\mathcal{A}$  is upper bounded by

$$\$AV((\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) \cdot (1 - p) + \mathbb{B}x_b \cdot p). \quad (2)$$

Since  $p \leq \gamma^{\tau^A} \leq \frac{\mathbb{A}c_a^A}{\mathbb{A}c_a^A + \mathbb{A}x_a}$ , we have

$$\$AV((\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) \cdot (1 - p) + \mathbb{B}x_b \cdot p) < \$AV(\mathbb{B}x_b - \mathbb{A}x_a).$$

- *Case 2:  $\mathcal{A}$  does not send any transaction containing  $pre_s$  before BobChain time  $T_1^B$ .* In this case, the honest Bob will send ping to  $A_{\text{defuse}}^A$  at BobChain time  $T_1^B$ . If  $A_{\text{defuse}}^A$  has not been activated at AliceChain time  $t^* \geq T_1^A$ , then, during AliceChain time  $(t^*, t^* + \tau^A]$ , any honest miner will activate  $A_{\text{burn}}^A$  if it wins a block. On the other hand, if  $A_{\text{defuse}}^A$  has been activated at AliceChain time  $t^* \geq T_1^A$ , the honest Bob will send  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated. Thus, at AliceChain time  $t^* \geq T_1^A$ ,  $pre_a$  and  $pre_b$  are both publicly known. Thus, during AliceChain time  $(t^*, t^* + \tau^A]$ , any honest miner will activate  $C_{\text{burn}}^A$  if it wins a block. By the same calculation as the previous case, since  $p \leq \gamma^{\tau^A} \leq \frac{\mathbb{A}c_a^A}{\mathbb{A}c_a^A + \mathbb{A}x_a}$ , we have  $\$AV((\mathbb{A}x_a - \mathbb{A}c_a^A + \mathbb{B}x_b) \cdot (1 - p) + \mathbb{B}x_b \cdot p) < \$AV(\mathbb{B}x_b - \mathbb{A}x_a)$ .

□

**Lemma 5.4** (Bob-miner coalition). *Suppose that the hash function  $H(\cdot)$  is a one-way function. Let  $\mathcal{B}$  be any coalition that consists of Bob and a subset of miners controlling at most  $\gamma$  fraction of mining power. Then, as long as  $\gamma^{\tau^B} \leq \frac{\mathbb{B}c_b^B}{\mathbb{B}c_b^B + \mathbb{B}x_b}$ , for any PPT coalition strategy  $S_{\mathcal{B}}$ , except with negligible probability, it must be*

$$\text{util}^{\mathcal{B}}(S_{\mathcal{B}}, HS_{-\mathcal{B}}) \leq \text{util}^{\mathcal{B}}(HS_{\mathcal{B}}, HS_{-\mathcal{B}}),$$

where  $HS_{-\mathcal{B}}$  denotes the honest strategy for everyone not in  $\mathcal{B}$ .

*Proof.* Recall that the utility of  $\mathcal{B}$  is  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b) > 0$  under an honest execution. Now, suppose  $\mathcal{B}$  may deviate from the protocol. We may assume that the coalition does not post any new smart contract on the fly and deposit money into it — if it did so, it cannot recover more than its deposit since any player not in  $\mathcal{B}$  will not invoke the smart contract. We analyze the two possible cases depending on which phase Alice enters.

**Alice enters the abort phase.** If  $\text{CONTRACT}^A$  never becomes active, the net profit of  $\mathcal{B}$  from AliceChain is at most zero. Now, assume  $\text{CONTRACT}^A$  is active. When Alice enters the abort phase, she never sends any transaction containing  $pre_s$ . Ignoring the negligible probability that  $\mathcal{B}$  finds  $pre_s$  by itself,  $P_{\text{default}}^A$  can never be activated, which means  $\text{Normal}^A$  never happens. According to Table 2, if  $\text{Normal}^A$  does not happen, the net profit of  $\mathcal{B}$  from AliceChain is at most zero. On the other hand, because  $\mathbb{B}x_b > \mathbb{B}c_b^B$ , the net profit of  $\mathcal{B}$  from BobChain is at most zero, no matter  $\text{CONTRACT}^B$  is active or not.

To sum up, except with negligible probability, the utility of  $\mathcal{B}$  is at most zero, which is less than the honest case.

**Alice enters the execution phase.** By Lemma 5.2, the utility of  $\mathcal{B}$  can be more than the honest case only if  $\text{Refund}^B$  and  $\text{Normal}^A$  both happen, so we assume it is the case. Because Alice enters the execution phase, both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  must be active. Thus,  $\text{Refund}^B$  happens only if  $P_{\text{refund}}^B$  is activated. When Alice enters the execution phase, she never sends ping to  $P_{\text{refund}}^B$ , so  $P_{\text{refund}}^B$  must be activated by  $pre_b$  sent by Bob. Therefore, we may assume that  $P_{\text{refund}}^B$  is activated at BobChain time  $t^* \geq T_1^B$ , and  $pre_b$  is publicly known after BobChain time  $t^*$ . If Alice enters the execution, Bob must have sent  $pre_c$  before BobChain time  $T_0$ . Moreover, Alice sends  $pre_s$  to  $P_{\text{default}}^B$  at BobChain time  $T_0$  and  $T_0 < T_1^B$ . Therefore,  $pre_s$ ,  $pre_b$  and  $pre_c$  are all publicly known at BobChain time  $t^*$ . Thus, during BobChain time  $(t^*, t^* + \tau^B]$ , any honest miner will activate  $C_{\text{burn}}^B$  if it wins a block. We say  $\mathcal{B}$  loses the race if a non-colluding miner mines a new block during BobChain time  $(t^*, t^* + \tau^B]$ . Otherwise, we say  $\mathcal{B}$  wins the race. If  $\mathcal{B}$  loses the race, it gets nothing from  $C_{\text{refund}}^B$  or  $C_{\text{burn}}^B$ , and its utility is at most  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B)$  which can be achieved by  $P_{\text{default}}^A$ . Else if  $\mathcal{B}$  wins the race, then its utility is at most  $\$BV(\mathbb{A}x_a)$  which can be achieved by activating  $P_{\text{refund}}^B$ ,  $C_{\text{refund}}^B$  and  $P_{\text{default}}^A$ . Since  $p \leq \gamma^{\tau^B} \leq \frac{\mathbb{B}c_b^B}{\mathbb{B}c_b^B + \mathbb{B}x_b}$ , we have

$$\$BV((\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B) \cdot (1 - p) + \mathbb{A}x_a \cdot p) < \$BV(\mathbb{A}x_a - \mathbb{B}x_b).$$

□

**Proof of Theorem 5.1.** Now, we are ready to prove Theorem 5.1. In Lemma 5.3 and Lemma 5.4, we show that the atomic swap protocol satisfies  $\gamma$ -CSP-fairness when the coalition consists of Alice or Bob, and possibly with some miners. Because we assume that Alice and Bob are not in the same coalition, it remains to show  $\gamma$ -CSP-fairness when the coalition  $\mathcal{C}$  consists only of miners controlling at most  $\gamma$  fraction of the mining power.

Henceforth, we assume Alice and Bob are both honest. It is clear from the protocol that the honest Alice and honest Bob always make the same decision whether to enter the execution phase or abort phase.

Next, when  $\mathcal{C}$  follows the protocol, its utility is always zero. Suppose  $\mathcal{C}$  may deviate from the protocol. Notice that the utility of  $\mathcal{C}$  can be positive only when  $A_{\text{burn}}^A$ ,  $B_{\text{burn}}^B$ ,  $C_{\text{burn}}^B$  or  $C_{\text{burn}}^A$  is activated. Because Bob only sends  $\text{pre}_c$  when  $B_{\text{defuse}}^B$  has been activated, ignoring the negligible probability that  $\mathcal{C}$  finds  $\text{pre}_c$  by itself,  $B_{\text{burn}}^B$  can never be activated. In the following, we will show that  $A_{\text{burn}}^A$ ,  $C_{\text{burn}}^B$  and  $C_{\text{burn}}^A$  are never activated except with negligible probability. There are two possible cases.

- *Case 1: both Alice and Bob enter the execution phase.* In this case, Alice always sends  $\text{pre}_s$  to  $P_{\text{default}}^B$ , and she never sends any transaction containing  $\text{pre}_a$ . Ignoring the negligible probability that  $\mathcal{C}$  finds  $\text{pre}_a$  by itself,  $C_{\text{burn}}^A$  can never be activated, and  $A_{\text{burn}}^A$  can only be activated by  $\text{pre}_b$ . Moreover, Alice always sends  $\text{pre}_s$  to  $P_{\text{default}}^B$  at latest at BobChain time  $T_0$ , and thus Bob will not post any transaction containing  $\text{pre}_b$ . Ignoring the negligible probability that  $\mathcal{C}$  finds  $\text{pre}_b$  by itself,  $A_{\text{burn}}^A$  and  $C_{\text{burn}}^B$  can never be activated. To sum up, except the negligible probability, the utility of  $\mathcal{C}$  is at most zero, which is the same as the honest case.
- *Case 2: both Alice and Bob enter the abort phase.* In this case, Bob always sends  $\text{ping}$  to  $P_{\text{refund}}^A$  and Alice always sends  $\text{ping}$  to  $P_{\text{refund}}^B$ . Thus, Bob never sends any transaction containing  $\text{pre}_b$ , and Alice never sends any transaction containing  $\text{pre}_a$ . Ignoring the negligible probability that  $\mathcal{C}$  finds  $\text{pre}_b$  or  $\text{pre}_a$  by itself,  $A_{\text{burn}}^A$ ,  $C_{\text{burn}}^B$  and  $C_{\text{burn}}^A$  cannot be activated by  $(\text{pre}_a, \text{pre}_b)$ . Thus, except with negligible probability, the utility of  $\mathcal{C}$  is at most zero, which is the same as the honest case.

## 5.2 Achieving Bounded Maximin Fairness

Henceforth, let  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$  denote the maximum fraction of mining power controlled by the set of externally incentivized players, and let  $\$E$  be an upper bound on any individual or coalition's valuation of the total possible external incentive.

**Proof Roadmap.** Conceptually, because Alice and Bob both put the collateral on both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$ , none of them wants to trigger any of the bombs ( $B_{\text{burn}}^B$ ,  $C_{\text{burn}}^B$ ,  $A_{\text{burn}}^A$  and  $C_{\text{burn}}^A$ ). In Section 5.2.1, we define a set of “bad events” that leads to the activation of the bomb. The bad events are defined such that if any of the bad events is about to happen, it must be that a strategic individual or coalition is about to send a transaction that does not follow from the honest protocol. In Section 5.2.2, we show that whenever the Bob-miner coalition  $\mathcal{B}$  is about to send a message that makes a bad event happen,  $\mathcal{B}$ 's expected utility can be strictly improved if  $\mathcal{B}$  simply stops sending any messages (including the message that is about to trigger the bad event) from that moment on (Lemma 5.7). Hence, any strategy that makes the bad event happen is a blatantly irrational strategy for  $\mathcal{B}$ , so a rational player would never make the bad events happen. Then, we show that as long as none of the bad events happens, the honest Alice's utility is never negative (Lemma 5.6 and Lemma 5.8). A similar argument can be made for the case of an Alice-miner coalition (see Section 5.2.3). Finally, in Section 5.2.4, we combine all the arguments above, and prove that the atomic swap protocol achieves bounded maximin fairness.

### 5.2.1 Irrational Strategies

In this section, we will define a family of PPT strategies denoted  $\overline{\mathcal{R}}$ , and we will show given any strategy  $S \in \overline{\mathcal{R}}$ , we can give a simple modification of  $S$ , resulting in a new PPT strategy which

makes the externally incentivized coalition better off — in this sense, the strategy space  $\overline{\mathcal{R}}$  is blatantly irrational.

**Terminologies.** Consider any user-miner coalition, and we define the following terminologies.

- We say an activation point  $X$  is *guaranteed to be activated* at some time  $T$ , iff either  $X$  was already activated before  $T$ , or a colluding miner has been chosen as the winning miner at time  $T$ , and it activates  $X$  in the new block it mines.
- We say a smart contract is *guaranteed to be active* at some time  $T$ , iff either the contract was already active before  $T$ , or a colluding miner has been chosen as the winning miner at time  $T$ , and it includes Alice's and Bob's deposit transactions for the contract in the new block it mines.
- We say Alice is *guaranteed to withdraw her deposit from a contract* at some time  $T$ , iff either Alice already withdrew her deposit from the contract before  $T$ , or all the following conditions hold.
  - The contract is not active before time  $T$ .
  - At time  $T$ , a colluding miner is chosen as the winning miner, and Alice's withdrawal transaction is included in the block at time  $T$ .

The case that Bob is guaranteed to withdraw his deposit from the contract is defined similarly.

**Irrational strategies.** The set  $\overline{\mathcal{R}}$  of irrational strategies for the externally incentivized Bob-miner coalition (including Bob alone)  $\mathcal{B}$  is the set of strategies such that with non-negligible probability, any of the following happens:

- E<sub>1</sub>:** Before  $B_{\text{defuse}}^B$  is guaranteed to be activated and before Bob is guaranteed to withdraw his deposit from  $\text{CONTRACT}^B$ , anyone in  $\mathcal{B}$  sends  $pre_c$  to  $(B_{\text{burn}}^B, P_{\text{default}}^B \text{ or } C_{\text{burn}}^B)$  and the deposit transaction to  $\text{CONTRACT}^B$ .
- E<sub>2</sub>:**  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  are guaranteed to be active before BobChain time  $T^B$ . Additionally, anyone in  $\mathcal{B}$  sends  $pre_c$  to  $(B_{\text{burn}}^B, P_{\text{default}}^B \text{ or } C_{\text{burn}}^B)$  before BobChain time  $T_0^B$ , and sends  $pre_b$  to  $(P_{\text{refund}}^B, C_{\text{burn}}^B, A_{\text{burn}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^B \text{ or } C_{\text{refund}}^B)$  is guaranteed to be activated.
- E<sub>3</sub>:** Before  $A_{\text{defuse}}^A$  is guaranteed to be activated and before Bob is guaranteed to withdraw his deposit from  $\text{CONTRACT}^A$ , anyone in  $\mathcal{B}$  sends  $pre_b$  to  $(P_{\text{refund}}^B, C_{\text{burn}}^B, A_{\text{burn}}^A \text{ or } C_{\text{burn}}^A)$  and the deposit transaction to  $\text{CONTRACT}^A$ .
- E<sub>4</sub>:** Alice enters the abort phase, and Bob does not send  $\text{ping}$  to  $P_{\text{refund}}^A$  before AliceChain time  $T_1^A$ . Additionally, at AliceChain time  $T_1^A$  or later, anyone in  $\mathcal{B}$  sends  $pre_b$  to  $(P_{\text{refund}}^B, C_{\text{burn}}^B, A_{\text{burn}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^A \text{ or } C_{\text{refund}}^A)$  is guaranteed to be activated.

The set  $\overline{\mathcal{R}}$  of irrational strategies for the externally incentivized Alice-miner coalition (including Alice alone)  $\mathcal{A}$  is the set of strategies such that with non-negligible probability, any of the following happens:

- E<sub>5</sub>:** Bob sends  $pre_c$  to  $P_{\text{default}}^B$  before BobChain time  $T_0^B$ , and Alice does not send  $pre_s$  to  $P_{\text{default}}^B$  until BobChain time  $T_1^B$ . However, at BobChain time  $T_1^B$  or later, anyone in  $\mathcal{A}$  sends  $pre_s$  to  $(P_{\text{default}}^B, C_{\text{burn}}^B, P_{\text{default}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^B \text{ or } C_{\text{refund}}^B)$  is guaranteed to be activated.

**E<sub>6</sub>:** Before  $A_{\text{defuse}}^A$  is guaranteed to be activated and before Alice is guaranteed to withdraw her deposit from  $\text{CONTRACT}^A$ , anyone in  $\mathcal{A}$  sends  $\text{pre}_a$  to  $(A_{\text{burn}}^A, P_{\text{refund}}^A \text{ or } C_{\text{burn}}^A)$  and the deposit transaction to  $\text{CONTRACT}^A$ .

**E<sub>7</sub>:** Anyone in  $\mathcal{A}$  sends  $\text{pre}_s$  to  $(P_{\text{default}}^B, C_{\text{burn}}^B, P_{\text{default}}^A \text{ or } C_{\text{burn}}^A)$  and  $\text{pre}_a$  to  $(A_{\text{burn}}^A, P_{\text{refund}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^A \text{ or } C_{\text{refund}}^A)$  is guaranteed to be activated.

**E<sub>8</sub>:** Any one of the conditions holds.

- Bob enters the execution phase, and Alice does not send  $\text{pre}_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ . Additionally, at AliceChain time  $T_1^A$  or later, anyone in  $\mathcal{A}$  sends  $\text{pre}_a$  to  $(A_{\text{burn}}^A, P_{\text{refund}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^A \text{ or } C_{\text{refund}}^A)$  is guaranteed to be activated.
- Bob enters the abort phase, and Alice does not send  $\text{ping}$  to  $P_{\text{refund}}^B$  before AliceChain time  $T_1^A$ . Additionally, at AliceChain time  $T_1^A$  or later, anyone in  $\mathcal{A}$  sends  $\text{pre}_a$  to  $(A_{\text{burn}}^A, P_{\text{refund}}^A \text{ or } C_{\text{burn}}^A)$  before  $(P_{\text{default}}^A \text{ or } C_{\text{refund}}^A)$  is guaranteed to be activated.

The following lemma specifies the upper bounds of the utility of the externally incentivized players.

**Lemma 5.5.** *Suppose the coalition  $\mathcal{B}$  consists of Bob and possibly some miners. Let  $\$E$  be an upper bound on  $\mathcal{B}$ 's valuation of the total possible external incentive. If  $\mathbb{B}c_b^B > \mathbb{B}\epsilon^B$ ,  $\mathbb{A}c_b^A > \mathbb{A}\epsilon^A$ , then, the following statements hold.*

- The utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a) + \$E$ .
- If Burn<sup>B</sup> is activated by an honest miner  $\notin \mathcal{B}$ , the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B) + \$E$ .
- If Burn<sup>A</sup> is activated by an honest miner  $\notin \mathcal{B}$ , the utility of  $\mathcal{B}$  is at most  $\$BV(-\mathbb{A}c_b^A) + \$E$ .

Similarly, suppose the coalition  $\mathcal{A}$  consists of Alice and possibly some miners. Let  $\$E$  be an upper bound on  $\mathcal{A}$ 's valuation of the total possible external incentive. If  $\mathbb{B}c_a^B > \mathbb{B}\epsilon^B$ ,  $\mathbb{A}c_a^A > \mathbb{A}\epsilon^A$ , then, the following statements hold.

- The utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b) + \$E$ .
- If Burn<sup>B</sup> is activated by an honest miner  $\notin \mathcal{A}$ , the utility of  $\mathcal{A}$  is at most  $\$AV(-\mathbb{B}c_a^B) + \$E$ .
- If Burn<sup>A</sup> is activated by an honest miner  $\notin \mathcal{A}$ , the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) + \$E$ .

*Proof.* The maximal possible utilities under different events are summarized in Table 1 and Table 2. The lemma directly follows from the calculation of the utilities in the table.  $\square$

### 5.2.2 Against Externally Incentivized Bob-Miner Coalition

**Lemma 5.6.** *Let  $\mathcal{B}$  be the coalition consisting of Bob and miners controlling no more than  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Suppose Alice and at least  $1 - \alpha$  fraction of mining power are honest. For any PPT strategy by  $\mathcal{B}$ , except with negligible probability, as long as none of **E<sub>1</sub>**, **E<sub>2</sub>**, **E<sub>3</sub>**, **E<sub>4</sub>** happens, then, one and only one of the following statement holds.*

1.  $\text{Normal}^B$  and  $\text{Normal}^A$  happen in polynomial time.

2.  $\text{Normal}^B$  and  $\text{Refund}^A$  happen in polynomial time.

3.  $\text{Refund}^B$  and  $\text{Refund}^A$  happen in polynomial time.

*Proof.* First, we are going to show that one of  $\text{Normal}^B$  and  $\text{Refund}^B$  will happen. There are two cases.

- *Case 1: Alice enters the execution phase.* We will show that  $P_{\text{default}}^B$  must be activated in polynomial time except with negligible probability. Once  $P_{\text{default}}^B$  is activated,  $\text{Normal}^B$  happens. In the execution phase, Alice always sends  $pre_s$  to  $P_{\text{default}}^B$ . If  $B_{\text{burn}}^B$ ,  $P_{\text{refund}}^B$  and  $C_{\text{burn}}^B$  are not activated, the honest miners will include Alice's transaction,  $pre_s$  to  $P_{\text{default}}^B$ , once they mine a block. Thus, it suffices to show that  $B_{\text{burn}}^B$ ,  $P_{\text{refund}}^B$ ,  $C_{\text{burn}}^B$  (all activation points that  $P_{\text{default}}^B$  is mutually exclusive with) cannot be activated except with negligible probability. First, because  $\mathbf{E}_1$  does not happen,  $\mathcal{B}$  never sends  $pre_c$  before  $B_{\text{defuse}}^B$  is guaranteed to be activated. Thus,  $B_{\text{burn}}^B$  can never be activated.

Next, because Alice enters the execution phase,  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  must be active, and Bob already sent  $pre_c$  to  $P_{\text{default}}^B$  before BobChain time  $T_0^B$ . Because  $\mathbf{E}_2$  does not happen, no one in  $\mathcal{B}$  sends  $pre_b$  to  $C_{\text{burn}}^B$  before  $P_{\text{default}}^B$  or  $C_{\text{refund}}^B$  is guaranteed to be activated. Thus,  $C_{\text{burn}}^B$  cannot be activated.

It remains to show that  $P_{\text{refund}}^B$  cannot be activated. In the execution phase, Alice never sends ping to  $P_{\text{refund}}^B$ , so  $P_{\text{refund}}^B$  can be activated only if Bob sends  $pre_b$  to  $P_{\text{refund}}^B$ . Because  $\mathbf{E}_2$  does not happen, Bob never sends  $pre_b$  to  $P_{\text{refund}}^B$  before  $P_{\text{default}}^B$  or  $C_{\text{refund}}^B$  is guaranteed to be activated. If  $P_{\text{default}}^B$  is guaranteed to be activated,  $P_{\text{refund}}^B$  cannot be activated as they are mutually exclusive. On the other hand, if  $C_{\text{refund}}^B$  is guaranteed to be activated,  $P_{\text{refund}}^B$  must have been activated  $\tau^B \geq 1$  BobChain time before. Thus, by the time  $P_{\text{refund}}^B$  is activated,  $C_{\text{refund}}^B$  has not been guaranteed to be activated. However, and Bob never sends  $pre_b$  to  $P_{\text{refund}}^B$  before  $C_{\text{refund}}^B$  is guaranteed to be activated. Therefore,  $P_{\text{refund}}^B$  can never be activated.

- *Case 2: Alice enters the abort phase.* In this case, Alice will send the withdrawal transaction to  $\text{CONTRACT}^B$ , and ping to  $P_{\text{refund}}^B$ . When  $\text{CONTRACT}^B$  has not been active yet, the honest miner will include Alice's withdrawal transactions once they mine a block. Thus, except with negligible probability, in polynomial time, either  $\text{CONTRACT}^B$  becomes active, or Alice successfully withdraws her deposit from  $\text{CONTRACT}^B$ . If Alice withdraws her deposit from  $\text{CONTRACT}^B$ ,  $\text{Refund}^B$  happens.

Henceforth, we assume  $\text{CONTRACT}^B$  becomes active. Because  $\mathbf{E}_1$  does not happen,  $\mathcal{B}$  never sends  $pre_c$  before  $B_{\text{defuse}}^B$  is guaranteed to be activated. Thus,  $B_{\text{burn}}^B$  can never be activated. Then, notice that Alice never sends  $pre_s$  when she enters the abort phase. Ignoring the negligible probability that  $\mathcal{B}$  finds  $pre_s$  by itself,  $C_{\text{burn}}^B$  can never be activated.

Thus, since  $1 - \alpha$  fraction of mining power is honest, either  $P_{\text{default}}^B$  or  $P_{\text{refund}}^B$  will be activated in polynomial time. If  $P_{\text{default}}^B$  is activated,  $\text{Normal}^B$  happens. If  $P_{\text{refund}}^B$  is activated, Alice will send ping to  $C_{\text{refund}}^B$  when  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated. Again, the honest miner will include Alice's transaction ping to  $C_{\text{refund}}^B$ , once they mine a block. Thus,  $\text{Refund}^B$  will happen in polynomial time.

Next, we are going to show that one of  $\text{Normal}^A$  and  $\text{Refund}^A$  will happen. There are two cases.

- *Case 1: Alice enters the execution phase.* As we have shown, when Alice enters the execution phase,  $P_{\text{default}}^B$  must be activated in polynomial time. Then, Alice will send ping to  $P_{\text{default}}^A$  as soon as  $P_{\text{default}}^B$  is activated. Because  $\mathbf{E}_3$  does not happen and Alice never sends  $pre_a$  when

in the execution phase,  $A_{\text{burn}}^A$  cannot be activated. Ignoring the negligible probability that  $\mathcal{B}$  finds  $\text{pre}_a$  by itself,  $C_{\text{burn}}^A$  cannot be activated. Because  $1 - \alpha$  fraction of mining power is honest, either  $P_{\text{default}}^A$  or  $P_{\text{refund}}^A$  will be activated in polynomial time. If  $P_{\text{default}}^A$  is activated,  $\text{Normal}^A$  happens. If  $P_{\text{refund}}^A$  is activated, Alice will send ping to  $C_{\text{refund}}^A$ . Again, the honest miner will include a transaction ping to  $C_{\text{refund}}^A$ , once they mine a block. Thus,  $C_{\text{refund}}^A$  will be activated in polynomial time, and  $\text{Refund}^A$  happens.

- *Case 2: Alice enters the abort phase.* In this case, Alice will send the withdrawal transaction to  $\text{CONTRACT}^A$ , and ping to  $A_{\text{defuse}}^A$ . If  $\text{CONTRACT}^A$  has not been active, the honest miner will include Alice's withdrawal transactions once they mine a block. Thus, in polynomial time, either  $\text{CONTRACT}^A$  becomes active, or Alice successfully withdraws her deposit from  $\text{CONTRACT}^A$ , so  $\text{Refund}^A$  happens.

Henceforth, we assume  $\text{CONTRACT}^A$  becomes active in polynomial time. Because  $\mathbf{E}_3$  does not happen,  $\mathcal{B}$  never sends  $\text{pre}_b$  before  $A_{\text{defuse}}^A$  is guaranteed to be activated. As  $A_{\text{defuse}}^A$  and  $A_{\text{burn}}^A$  are mutually exclusive,  $A_{\text{burn}}^A$  can not be activated via  $\text{pre}_b$ . Since Alice sends  $\text{pre}_a$  to  $P_{\text{refund}}^A$  only after  $A_{\text{defuse}}^A$  is activated, 1)  $A_{\text{burn}}^A$  can not be activated via  $\text{pre}_a$ , and 2)  $C_{\text{burn}}^A$  can not happen before  $A_{\text{defuse}}^A$  is activated. If either  $P_{\text{default}}^A$  or  $C_{\text{refund}}^A$  are activated,  $\text{Normal}^A$  or  $\text{Refund}^A$  happens. Otherwise, the honest miner will include the transaction ping to  $A_{\text{defuse}}^A$  once they mine a block. Because  $1 - \alpha$  fraction of mining power is honest,  $A_{\text{defuse}}^A$  will be activated in polynomial time.

When Alice is in abort phase, either Bob sends ping to  $P_{\text{refund}}^A$  before AliceChain time  $T_1^A$  or Alice sends  $\text{pre}_a$  to  $P_{\text{refund}}^A$  when  $A_{\text{defuse}}^A$  is activated. If Bob sends ping to  $P_{\text{refund}}^A$  before AliceChain time  $T_1^A$ , Alice never sends  $\text{pre}_a$ , and thus  $C_{\text{burn}}^A$  can not be activated. When Alice enters the abort phase, she never sends  $\text{pre}_s$ , and thus  $P_{\text{default}}^A$  can not be activated. As we showed before,  $A_{\text{burn}}^A$  can not be activated either, and thus the honest miner will include Bob's transaction, ping to  $P_{\text{refund}}^A$ , once they mine a block, so  $P_{\text{refund}}^A$  will be activated in polynomial time.

On the other hand, suppose Bob does not send ping to  $P_{\text{refund}}^A$  before AliceChain time  $T_1^A$ , so Alice sends  $\text{pre}_a$  to  $P_{\text{refund}}^A$ . Because  $\mathbf{E}_4$  does not happen,  $\mathcal{B}$  never sends  $\text{pre}_b$  before  $P_{\text{default}}^A$  or  $C_{\text{refund}}^A$  is guaranteed to be activated. If  $P_{\text{default}}^A$  is activated,  $\text{Normal}^A$  happens, if  $C_{\text{refund}}^A$  is activated,  $P_{\text{refund}}^A$  must have been activated previously, and thus  $\text{Refund}^A$  happens. Otherwise, recall that Alice never sends  $\text{pre}_s$  in the abort phase. Without  $\text{pre}_s$  and  $\text{pre}_b$ ,  $P_{\text{default}}^A$  and  $C_{\text{burn}}^A$  cannot be activated, and as  $A_{\text{burn}}^A$  can not be activated either as shown before, the honest miner will include Alice's transaction,  $\text{pre}_a$  to  $P_{\text{refund}}^A$ , once they mine a block. Thus,  $P_{\text{refund}}^A$  will again be activated in polynomial time.

Then, Alice will send ping to  $C_{\text{refund}}^A$  when  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated. Again, the honest miner will include Alice's transaction, ping to  $C_{\text{refund}}^A$ , once they mine a block. Thus,  $\text{Refund}^A$  will happen in polynomial time.

So far, we have shown one of  $\text{Normal}^B$  and  $\text{Refund}^B$  will happen and one of  $\text{Normal}^A$  and  $\text{Refund}^A$  will happen. To prove the lemma statement, it suffices to show that  $\text{Refund}^B$  and  $\text{Normal}^A$  never happen simultaneously. For the sake of reaching a contradiction, suppose  $\text{Refund}^B$  and  $\text{Normal}^A$  happen. Event  $\text{Normal}^A$  happens implies  $P_{\text{default}}^A$  is activated, and  $P_{\text{default}}^A$  is activated only when Bob sends  $\text{pre}_s$  or Alice sends ping. The honest Alice only sends ping to  $P_{\text{default}}^A$  when  $P_{\text{default}}^B$  is activated, which implies  $\text{Normal}^B$  happens. Henceforth, we assume  $P_{\text{default}}^A$  is activated when by Bob's  $\text{pre}_s$ . Ignoring the negligible probability that  $\mathcal{B}$  can find  $\text{pre}_s$  by itself, Alice must enter the execution phase and send  $\text{pre}_s$  to  $P_{\text{default}}^B$  at BobChain time  $T_0^B$ . On the other hand, event

$\text{Refund}^B$  happens implies either  $(P_{\text{refund}}^B + C_{\text{refund}}^B)$  are activated or Bob withdraws his deposit from  $\text{CONTRACT}^B$ . Because Alice enters the execution phase, both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  must also enter the execution, and thus we exclude that Bob withdraws his deposit from  $\text{CONTRACT}^B$ . Thus, event  $\text{Refund}^B$  happens only when  $(P_{\text{refund}}^B + C_{\text{refund}}^B)$  are activated. Moreover, because Alice enters the execution phase, Bob must have sent  $\text{pre}_c$  to  $P_{\text{default}}^B$ . Therefore, to activate  $P_{\text{refund}}^B$ , Bob has to send  $\text{pre}_b$  before  $C_{\text{refund}}^B$  is guaranteed to be activated, which implies event  $\mathbf{E}_2$  happens.  $\square$

**Lemma 5.7** (Blatant irrationality of  $\bar{\mathcal{R}}$  for Bob-miner coalition). *Suppose that the parameter constraints in Figure 2 hold and that the coalition  $\mathcal{B}$  consists of Bob and miners controlling no more than  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Given any PPT strategy  $S_{\mathcal{B}} \in \bar{\mathcal{R}}$  for some (externally incentivized) coalition  $\mathcal{B}$ , there is a PPT strategy  $\hat{S}_{\mathcal{B}}$  such that*

$$\text{util}^{\mathcal{B}}(\hat{S}_{\mathcal{B}}, HS_{-\mathcal{B}}) > \text{util}^{\mathcal{B}}(S_{\mathcal{B}}, HS_{-\mathcal{B}}).$$

*Proof.* Suppose the coalition  $\mathcal{B}$  adopts a strategy in which  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  or  $\mathbf{E}_4$  happens with non-negligible probability. We can construct a new strategy for  $\mathcal{B}$  with strictly better expected utility. Specifically, consider a modified PPT strategy denoted  $\hat{S}_{\mathcal{B}}$ : whenever by the original strategy, the first of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  or  $\mathbf{E}_4$  is about to happen,  $\mathcal{B}$  simply stops sending any messages (including the message that is about to trigger  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  or  $\mathbf{E}_4$ ) to the contract from that moment on.

By definition, when  $\mathbf{E}_1$  or  $\mathbf{E}_3$  happens, it must due to a message sent by  $\mathcal{B}$ . According to the protocol, honest Alice always enters the execution or the abort phase no later than BobChain time  $T_0^B$ . Because AliceChain time  $T_1^A$  is later than BobChain time  $T_0^B$  by the choice of the parameters, when  $\mathbf{E}_4$  happens, it must also due to a message sent by  $\mathcal{B}$ . Next, if  $\text{CONTRACT}^A$  has not been guaranteed to be active,  $A_{\text{defuse}}^A$  cannot be guaranteed to be activated. Similarly, if  $\text{CONTRACT}^B$  has not been guaranteed to be active,  $B_{\text{defuse}}^B$  cannot be guaranteed to be activated. Thus, if anyone in  $\mathcal{B}$  sends  $\text{pre}_c$  to  $(B_{\text{burn}}^B, P_{\text{default}}^B \text{ or } C_{\text{burn}}^B)$  and sends  $\text{pre}_b$  to  $(P_{\text{refund}}^B, C_{\text{burn}}^B, A_{\text{burn}}^A \text{ or } C_{\text{burn}}^A)$  before  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  both are guaranteed to be active, one of  $\mathbf{E}_1$  and  $\mathbf{E}_3$  must happen. Consequently, if  $\mathbf{E}_1$  and  $\mathbf{E}_3$  do not happen while  $\mathbf{E}_2$  happens, it must due to a message sent by  $\mathcal{B}$ . Thus, as long as  $\mathcal{B}$  stops sending any messages before the first of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  or  $\mathbf{E}_4$  is about to happen, none of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  and  $\mathbf{E}_4$  can happen in the future. By Lemma 5.6, in polynomial time, one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Normal}^B + \text{Refund}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$  must happen. By direct calculation, we have the following table. Among  $(\text{Normal}^B + \text{Normal}^A)$ ,

	$\text{Normal}^B$	$\text{Refund}^B$	$\text{Burn}^B$
$\text{Normal}^A$	$-\ddot{\mathbb{B}}x_b + \ddot{\mathbb{A}}x_a$	$\ddot{\mathbb{A}}x_a$	$-\ddot{\mathbb{B}}x_b - \ddot{\mathbb{B}}c_b^B + \ddot{\mathbb{A}}x_a$
$\text{Refund}^A$	$-\ddot{\mathbb{B}}x_b$	0	$-\ddot{\mathbb{B}}x_b - \ddot{\mathbb{B}}c_b^B$
$\text{Burn}^A$	$-\ddot{\mathbb{B}}x_b - \ddot{\mathbb{A}}c_b^A$	$-\ddot{\mathbb{A}}c_b^A$	$-\ddot{\mathbb{B}}x_b - \ddot{\mathbb{B}}c_b^B - \ddot{\mathbb{A}}c_b^A$

Table 3: Bob's net profit under all possible events.

$(\text{Normal}^B + \text{Refund}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$ ,  $\mathcal{B}$ 's utility is at least  $\$BV(-\ddot{\mathbb{B}}x_b)$ . In other words, we have  $\text{util}^{\mathcal{B}}(\hat{S}_{\mathcal{B}}, HS_{-\mathcal{B}}) \geq \$BV(-\ddot{\mathbb{B}}x_b)$ . Thus, to show  $\text{util}^{\mathcal{B}}(\hat{S}_{\mathcal{B}}, HS_{-\mathcal{B}}) \geq \text{util}^{\mathcal{B}}(S_{\mathcal{B}}, HS_{-\mathcal{B}})$ , we only need to show  $\text{util}^{\mathcal{B}}(S_{\mathcal{B}}, HS_{-\mathcal{B}}) < \$BV(-\ddot{\mathbb{B}}x_b)$ .

We consider four cases, depending on whether  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3$  or  $\mathbf{E}_4$  happens first in the original strategy  $S_{\mathcal{B}}$ .

**Event  $\mathbf{E}_1$  happens first.** Consider some strategy  $S_1 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_1$  happens. When  $\mathbf{E}_1$  happens, because Bob is not guaranteed to withdraw his deposit from

$\text{CONTRACT}^B$ , the  $1 - \alpha$  fraction of honest miners would send the deposit to  $\text{CONTRACT}^B$  and  $pre_c$  to  $B_{\text{burn}}^B$  (and potentially  $C_{\text{burn}}^B$  as well if  $pre_s$  and  $pre_b$  are available), if they are chosen to mine a block. By Lemma 5.5, if  $B_{\text{burn}}^B$  or  $C_{\text{burn}}^B$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B) + \$E$ . On the other hand, if neither  $B_{\text{burn}}^B$  nor  $C_{\text{burn}}^B$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a) + \$E$ . When  $\mathbf{E}_1$  happens, the probability that  $B_{\text{burn}}^B$  or  $C_{\text{burn}}^B$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, the utility of  $\mathcal{B}$  is at most

$$(1 - \alpha)(\$BV(\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B) + \$E) + \alpha(\$BV(\mathbb{A}x_a) + \$E) < \$BV(-\mathbb{B}x_b),$$

where the inequality arises from the fact that  $\$BV(\mathbb{B}c_b^B) > \frac{\$BV(\mathbb{A}x_a + \alpha\mathbb{B}x_b) + \$E}{1 - \alpha}$ .

**Event  $\mathbf{E}_2$  happens first.** Consider some strategy  $S_2 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_2$  happens. When  $\mathbf{E}_2$  happens, because both  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  become active before BobChain time  $T^B$  and Bob sends  $pre_c$  before BobChain time  $T_0^B$ , Alice enters the execution phase. In this case, Alice always sends  $pre_s$  to  $P_{\text{default}}^B$  at BobChain time  $T_0^B$ . However, Bob also sends  $pre_b$  before  $C_{\text{refund}}^B$  is guaranteed to be activated. Thus, the  $1 - \alpha$  fraction of honest miners would send  $(pre_s, pre_b, pre_c)$  to  $C_{\text{burn}}^B$  (and potentially send  $pre_c$  to  $B_{\text{burn}}^B$ ), if they are chosen to mine a block. By Lemma 5.5, if  $C_{\text{burn}}^B$  or  $B_{\text{burn}}^B$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b - \mathbb{B}c_b^B) + \$E$ . On the other hand, if neither  $C_{\text{burn}}^B$  nor  $B_{\text{burn}}^B$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a) + \$E$ . When  $\mathbf{E}_2$  happens, the probability that  $C_{\text{burn}}^B$  or  $B_{\text{burn}}^B$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, by the same calculation as the previous case, the utility of  $\mathcal{B}$  is strictly less than  $\$BV(-\mathbb{B}x_b)$ .

**Event  $\mathbf{E}_3$  happens first.** Consider some strategy  $S_3 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_3$  happens. When  $\mathbf{E}_3$  happens, because Bob is not guaranteed to withdraw his deposit from  $\text{CONTRACT}^A$ , the  $1 - \alpha$  fraction of honest miners would send the deposit to  $\text{CONTRACT}^A$  and  $pre_b$  to  $A_{\text{burn}}^A$  (and to  $C_{\text{burn}}^A$  if  $pre_a$  is available as well), if they are chosen to mine a block. By Lemma 5.5, if  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(-\mathbb{A}c_b^A) + \$E$ . On the other hand, if neither  $A_{\text{burn}}^A$  nor  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a) + \$E$ . When  $\mathbf{E}_3$  happens, the probability that  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, the utility of  $\mathcal{B}$  is at most

$$(1 - \alpha)(\$BV(-\mathbb{A}c_b^A) + \$E) + \alpha(\$BV(\mathbb{A}x_a) + \$E) < \$BV(-\mathbb{B}x_b),$$

where the inequality arises from the fact that  $\$BV(\mathbb{A}c_b^A) > \frac{\$BV(\mathbb{B}x_b + \alpha\mathbb{A}x_a) + \$E}{1 - \alpha}$ .

**Event  $\mathbf{E}_4$  happens first.** Consider some strategy  $S_4 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_4$  happens. When  $\mathbf{E}_4$  happens, because Alice enters the abort phase and Bob does not send  $\text{ping}$  to  $P_{\text{refund}}^A$  before AliceChain time  $T_1^A$ , Alice will send  $pre_a$  as soon as  $A_{\text{defuse}}^A$  is activated. However, Bob also sends  $pre_b$  before  $P_{\text{default}}^A$  or  $C_{\text{refund}}^A$  is guaranteed to be activated. If Bob sends  $pre_b$  before  $A_{\text{defuse}}^A$  is activated, the  $1 - \alpha$  fraction of honest miners would send  $pre_b$  to  $A_{\text{burn}}^A$ , if they are chosen to mine a block. If Bob sends  $pre_b$  after  $A_{\text{defuse}}^A$  is activated, the  $1 - \alpha$  fraction of honest miners would send  $(pre_a, pre_b)$  to  $C_{\text{burn}}^A$ , if they are chosen to mine a block. By Lemma 5.5, if  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(-\mathbb{A}c_b^A) + \$E$ . On the other hand, if neither  $A_{\text{burn}}^A$  nor  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{B}$  is at most  $\$BV(\mathbb{A}x_a) + \$E$ . When  $\mathbf{E}_4$  happens, the probability that  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, by the same calculation as the previous case, the utility of  $\mathcal{B}$  is strictly less than  $\$BV(-\mathbb{B}x_b)$ .

Finally, notice that the above analysis holds even if  $\mathcal{B}$  may post a new contract on the fly during the protocol execution, since all other players are honest and will not deposit money into the new contract.  $\square$

**Lemma 5.8** (Against Externally Incentivized Bob-Miner Coalition). *Suppose that the hash function  $H(\cdot)$  is a one-way function. Let  $\mathcal{C}$  be a group consisting of Alice and possibly any subset of the miners, and let  $\mathcal{B}$  be a disjoint coalition consisting of Bob and at most  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Suppose that  $\mathcal{C}$  does not have external incentives but  $\mathcal{B}$  may have up to  $\$E$  amount of external incentives. Let  $S_{\mathcal{B}}$  be an arbitrary PPT strategy of  $\mathcal{B}$  that is not in  $\bar{\mathcal{R}}$ . Then, there exists a negligible function  $\text{negl}(\cdot)$  such that except with negligible probability  $\text{negl}(\lambda)$ , it holds that*

$$\text{util}^{\mathcal{C}}(HS_{\mathcal{C}}, S_{\mathcal{B}}, HS_{\mathcal{D}}) \geq 0,$$

where  $\mathcal{D}$  denotes everyone else not in  $\mathcal{C} \cup \mathcal{B}$ .

*Proof.* By Lemma 5.7, any strategy that makes one of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3, \mathbf{E}_4$  happen is blatantly irrational. By direct calculation, we have the following table.

	Normal <sup>B</sup>	Refund <sup>B</sup>	Burn <sup>B</sup>
Normal <sup>A</sup>	$\$x_b - \Delta x_a$	$-\Delta x_a$	$-\$c_a^B - \Delta x_a$
Refund <sup>A</sup>	$\$x_b$	0	$-\$c_a^B$
Burn <sup>A</sup>	$\$x_b - \Delta x_a - \Delta c_a^A$	$-\Delta x_a - \Delta c_a^A$	$-\$c_a^B - \Delta x_a - \Delta c_a^A$

Table 4: Alice's net profit under all possible events.

By Lemma 5.6, if none of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3, \mathbf{E}_4$  happens, then one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Normal}^B + \text{Refund}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$  must happen. Because  $\$AV(\$x_b - \Delta x_a) > 0$ , except with some negligible probability, for all three possible cases Alice's utility is non-negative.  $\square$

### 5.2.3 Against Externally Incentivized Alice-Miner Coalition

**Lemma 5.9.** *Let  $\mathcal{A}$  be the coalition consisting of Alice and miners controlling no more than  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Suppose Bob and at least  $1 - \alpha$  fraction of mining power are honest. For any PPT strategy by  $\mathcal{A}$ , except with negligible probability, as long as none of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7, \mathbf{E}_8$  happens, then, one and only one of the following statement holds.*

1.  $\text{Normal}^B$  and  $\text{Normal}^A$  happen in polynomial time.
2.  $\text{Refund}^B$  and  $\text{Normal}^A$  happen in polynomial time.
3.  $\text{Refund}^B$  and  $\text{Refund}^A$  happen in polynomial time.

*Proof.* First, we are going to show that one of  $\text{Normal}^B$  and  $\text{Refund}^B$  will happen. There are two cases.

- *Case 1: Bob enters the execution phase.* In this case,  $\text{CONTRACT}^B$  and  $\text{CONTRACT}^A$  must be active, and Bob already sent  $\text{pre}_c$  to  $P_{\text{default}}^B$  before BobChain time  $T_0^B$ . There are two subcases.

- *Subcase 1: Alice sends  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ .* In this case, Bob never sends  $pre_b$ . Without  $pre_b$ ,  $C_{\text{burn}}^B$  can never be activated. As Bob is in the execution phase,  $B_{\text{defuse}}^B$  was activated and thus  $B_{\text{burn}}^B$  can never be activated. As Alice sent  $pre_s$  to  $P_{\text{default}}^B$  and Bob sent  $pre_c$  to  $P_{\text{default}}^B$ , either  $P_{\text{default}}^B$  is activated in polynomial time, or  $P_{\text{refund}}^B$  is activated instead. If  $P_{\text{default}}^B$  is activated,  $\text{Normal}^B$  happens. If  $P_{\text{refund}}^B$  is activated, Bob sends an empty message to  $C_{\text{refund}}^B$ , and thus  $C_{\text{refund}}^B$  is activated in polynomial time and  $\text{Refund}^B$  happens.
- *Subcase 2: Alice does not send  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ .* In this case, Bob will send ping to  $A_{\text{defuse}}^A$ . As  $\mathbf{E}_6$  does not happen, Alice never sends  $pre_a$  before  $A_{\text{defuse}}^A$  is guaranteed to be activated, and thus  $A_{\text{burn}}^A$  is never activated. Thus,  $A_{\text{defuse}}^A$  is activated in polynomial time and Bob sends  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated. Because  $\mathbf{E}_5$  does not happen, Alice never sends  $pre_s$  before  $P_{\text{default}}^B$  or  $C_{\text{refund}}^B$  is activated. Thus,  $C_{\text{burn}}^B$  can never be activated. As Bob is in the execution phase,  $B_{\text{burn}}^B$  can never be activated. Thus, either  $P_{\text{default}}^B$  is activated ( $\text{Normal}^B$  happens), or  $C_{\text{refund}}^B$  is activated (thus  $P_{\text{refund}}^B$  has been activated before and  $\text{Refund}^B$  happens), or the honest miner will include Bob’s transaction,  $pre_b$  to  $P_{\text{refund}}^B$ , once they mine a block. Then, when  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Bob will send ping to  $C_{\text{refund}}^B$  and thus  $\text{Refund}^B$  happens.
- *Case 2: Bob enters the abort phase.* In this case, Bob will send the withdrawal transaction to  $\text{CONTRACT}^B$ , and ping to  $A_{\text{defuse}}^A$ . If  $\text{CONTRACT}^B$  has not been active, the honest miner will include Bob’s withdrawal transactions once they mine a block. Thus, in polynomial time, either  $\text{CONTRACT}^B$  becomes active, or Bob successfully withdraws his deposit from  $\text{CONTRACT}^B$ , which implies  $\text{Refund}^B$  happens.

Henceforth, we assume  $\text{CONTRACT}^B$  becomes active in polynomial time. Because Bob is in the abort phase, Bob never sent  $pre_c$  and thus, up to negligible probability,  $B_{\text{burn}}^B$  and  $C_{\text{burn}}^B$  can not be activated. Because  $\mathbf{E}_6$  does not happen,  $A$  never sends  $pre_a$  before  $A_{\text{defuse}}^A$  is guaranteed to be activated. Thus,  $A_{\text{defuse}}^A$  is activated in polynomial time. Then, either Alice sends ping to  $P_{\text{refund}}^B$  before AliceChain time  $T_1^A$ , or Bob sends  $pre_b$  to  $P_{\text{refund}}^B$ . Either way, either  $P_{\text{default}}^B$  is activated (thus  $\text{Normal}^B$  happens), or  $P_{\text{refund}}^B$  is activated in polynomial time. When  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Bob will send ping to  $C_{\text{refund}}^B$ . Thus,  $C_{\text{refund}}^B$  will be activated, and so  $\text{Refund}^B$  happens.

Next, we are going to show that one of  $\text{Normal}^A$  and  $\text{Refund}^A$  will happen. There are two cases.

- *Case 1: Bob enters the execution phase.* First, because  $\mathbf{E}_6$  does not happen, Alice never sends  $pre_a$  before  $A_{\text{defuse}}^A$  is guaranteed to be activated, and as Bob sends  $pre_b$  only after  $A_{\text{defuse}}^A$  is activated, it follows that  $A_{\text{burn}}^A$  is never activated.

Suppose Alice sends  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ . Then, Bob never sends  $pre_b$ , and so  $C_{\text{burn}}^A$  can not be activated via  $(pre_a, pre_b)$ . Because  $\mathbf{E}_7$  does not happen,  $C_{\text{burn}}^A$  can not be activated via  $(pre_s, pre_a)$ . Thus,  $C_{\text{burn}}^A$  is never activated. As Bob sends  $pre_s$  to  $P_{\text{default}}^A$ , either  $\text{Normal}^A$  happens, or  $P_{\text{refund}}^A$  is activated. In the latter case, Bob sends ping to  $C_{\text{refund}}^A$ , and so  $\text{Refund}^A$  happens.

Henceforth, we assume Alice does not send  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ . As we have shown before, either  $\text{Normal}^B$  happens, or  $\text{Refund}^B$  happens. If  $\text{Normal}^B$  happens, Bob sends  $pre_s$  to  $P_{\text{default}}^A$ . If  $\text{Refund}^B$  happens, Bob will send ping to  $P_{\text{refund}}^A$  as soon as  $C_{\text{refund}}^B$  is activated. As  $\mathbf{E}_8$  does not happen,  $C_{\text{burn}}^A$  can not be activated using  $pre_a$  after AliceChain time  $T_1^A$ . As  $\mathbf{E}_6$  does not happen,  $C_{\text{burn}}^A$  can not be activated using  $pre_a$  before AliceChain

time  $T_1^A$  either. Thus,  $C_{\text{burn}}^A$  is never activated. Thus, in both cases (Normal<sup>B</sup> and Refund<sup>B</sup>), either  $P_{\text{default}}^A$  or  $P_{\text{refund}}^A$  will be activated in polynomial time. If  $P_{\text{default}}^A$  is activated, Normal<sup>A</sup> happens. If  $P_{\text{refund}}^A$  is activated, Bob will send ping to  $C_{\text{refund}}^A$ . The honest miners will include Bob's transaction, ping to  $C_{\text{refund}}^A$ , once they mine a block. Thus,  $C_{\text{refund}}^A$  will be activated in polynomial time, so Refund<sup>A</sup> happens.

- *Case 2: Bob enters the abort phase.* In this case, Bob will send the withdrawal transaction to CONTRACT<sup>A</sup>, ping to  $P_{\text{refund}}^A$  and ping to  $A_{\text{defuse}}^A$ . If CONTRACT<sup>A</sup> has not been active, the honest miner will include Bob's withdrawal transactions once they mine a block. Thus, in polynomial time, either CONTRACT<sup>A</sup> becomes active, or Bob successfully withdraws his deposit from CONTRACT<sup>A</sup>, which implies Refund<sup>A</sup> happens.

Henceforth, we assume CONTRACT<sup>A</sup> becomes active in polynomial time. Because  $\mathbf{E}_6$  does not happen,  $\mathcal{A}$  never sends  $\text{pre}_a$  before  $A_{\text{defuse}}^A$  is guaranteed to be activated. As Bob's sends ping to  $A_{\text{defuse}}^A$ ,  $A_{\text{defuse}}^A$  will be activated in polynomial time. If Alice sends ping to  $P_{\text{refund}}^B$  by AliceChain time  $T_1^A$ , Bob never sends  $\text{pre}_b$ , and thus  $C_{\text{burn}}^A$  can be activated only via ( $\text{pre}_s, \text{pre}_a$ ) (excluded by  $\mathbf{E}_7$ ). If Alice does not send ping to  $P_{\text{refund}}^B$  by AliceChain time  $T_1^A$ , because  $\mathbf{E}_8$  does not happen, Alice does not send  $\text{pre}_a$  before  $P_{\text{default}}^A$  or  $C_{\text{refund}}^A$  is guaranteed to be activated. Thus,  $C_{\text{burn}}^A$  can not be activated via  $\text{pre}_a$ . Thus,  $C_{\text{burn}}^A$  is never activated. As Bob sends ping to  $P_{\text{refund}}^A$ , or Alice's transaction, either  $P_{\text{default}}^A$  or  $P_{\text{refund}}^A$  will be activated in polynomial time. If  $P_{\text{default}}^A$  is activated, Normal<sup>A</sup> happens. If  $P_{\text{refund}}^A$  is activated, Bob will send ping to  $C_{\text{refund}}^A$ . Thus,  $C_{\text{refund}}^A$  will be activated in polynomial time, so Refund<sup>A</sup> happens.

So far, we have shown one of Normal<sup>B</sup> and Refund<sup>B</sup> will happen and one of Normal<sup>A</sup> and Refund<sup>A</sup> will happen. To prove the lemma statement, it suffices to show that Normal<sup>B</sup> and Refund<sup>A</sup> never happen simultaneously. For the sake of reaching a contradiction, suppose Normal<sup>B</sup> and Refund<sup>A</sup> happen. Event Normal<sup>B</sup> happens implies  $P_{\text{default}}^B$  is activated. If Bob enters the abort phase, Bob never sends  $\text{pre}_c$  to  $P_{\text{default}}^B$ . Ignoring the negligible probability that  $\mathcal{A}$  finds  $\text{pre}_c$  by itself and forges Bob's signature, Bob must enter the execution, which implies both CONTRACT<sup>B</sup> and CONTRACT<sup>A</sup> are active. On the other hand, event Refund<sup>A</sup> happens implies either ( $P_{\text{refund}}^A + C_{\text{refund}}^A$ ) are activated or Alice withdraws her deposit from CONTRACT<sup>A</sup>. As we have shown, CONTRACT<sup>A</sup> must be active, so Refund<sup>A</sup> happens implies ( $P_{\text{refund}}^A + C_{\text{refund}}^A$ ) are activated. In the execution phase, Bob will send ping to  $P_{\text{refund}}^A$  only if  $C_{\text{refund}}^A$  is activated, which is impossible given that  $P_{\text{default}}^B$  is activated. Thus,  $P_{\text{refund}}^A$  can be activated only if Alice sends  $\text{pre}_a$  to  $P_{\text{refund}}^A$ . Because  $C_{\text{refund}}^A$  can be activated when  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated, Alice must send  $\text{pre}_a$  to  $P_{\text{refund}}^A$  before  $C_{\text{refund}}^A$  is guaranteed to be activated. There are two subcases.

- *Subcase 1:  $\text{pre}_s$  is sent to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ .* Because  $P_{\text{refund}}^A$  must be activated after AliceChain time  $T_1^A$  which is later than BobChain time  $T_1^B$ ,  $\mathcal{A}$  must send  $\text{pre}_s$  to  $P_{\text{default}}^B$  and  $\text{pre}_a$  to  $P_{\text{refund}}^A$  before  $C_{\text{refund}}^A$  is guaranteed to be activated. Thus, depends on whether  $A_{\text{defuse}}^A$  is guaranteed to be activated, either  $\mathbf{E}_6$  or  $\mathbf{E}_7$  happens.
- *Subcase 2:  $\text{pre}_s$  has not been sent to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ .* Because Bob enters the execution phase, he will send  $\text{pre}_b$  to  $P_{\text{refund}}^A$  as soon as  $A_{\text{defuse}}^A$  is activated. Thus, depending on whether  $A_{\text{defuse}}^A$  is guaranteed to be activated, either  $\mathbf{E}_6$  or  $\mathbf{E}_8$  happens.

Because we assume none of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7, \mathbf{E}_8$  happens, either subcase leads to a contradiction.  $\square$

**Lemma 5.10** (Blatant irrationality of  $\bar{\mathcal{R}}$  for Alice-miner coalition). *Suppose that the parameter constraints in Figure 2 hold and that the coalition  $\mathcal{A}$  consists of Alice and miners controlling no*

more than  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Given any PPT strategy  $S_{\mathcal{A}} \in \bar{\mathcal{R}}$  for some (externally incentivized) coalition  $\mathcal{A}$ , there is a PPT strategy  $\hat{S}_{\mathcal{A}}$  such that

$$\text{util}^{\mathcal{A}}(\hat{S}_{\mathcal{A}}, HS_{-\mathcal{A}}) > \text{util}^{\mathcal{A}}(S_{\mathcal{A}}, HS_{-\mathcal{A}}).$$

*Proof.* Suppose the coalition  $\mathcal{A}$  adopts a strategy in which  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$  happens with non-negligible probability. We can construct a new strategy for  $\mathcal{A}$  with strictly better expected utility. Specifically, consider a modified PPT strategy denoted  $\hat{S}_{\mathcal{A}}$ : whenever by the original strategy, the first of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$  is about to happen,  $\mathcal{A}$  simply stops sending any messages (including the message that is about to trigger  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$ ) to the contract from that moment on.

By definition, when  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$  happens, it must be caused by a message sent by  $\mathcal{A}$ . Thus, as long as  $\mathcal{A}$  stops sending any messages before the first of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$  is about to happen, none of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  and  $\mathbf{E}_8$  can happen in the future. By Lemma 5.9, in polynomial time, one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Refund}^B + \text{Normal}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$  must happen. According to Table 4, among  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Refund}^B + \text{Normal}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$ ,  $\mathcal{A}$ 's utility is at least  $\$AV(-\mathbb{A}x_a)$ . In other words, we have  $\text{util}^{\mathcal{A}}(\hat{S}_{\mathcal{A}}, HS_{-\mathcal{A}}) \geq \$AV(-\mathbb{A}x_a)$ . Thus, to show  $\text{util}^{\mathcal{A}}(\hat{S}_{\mathcal{A}}, HS_{-\mathcal{A}}) \geq \text{util}^{\mathcal{A}}(S_{\mathcal{A}}, HS_{-\mathcal{A}})$ , we only need to show  $\text{util}^{\mathcal{A}}(S_{\mathcal{A}}, HS_{-\mathcal{A}}) < \$AV(-\mathbb{A}x_a)$ .

We consider four cases, depending on whether  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7$  or  $\mathbf{E}_8$  happens first in the original strategy  $S_{\mathcal{A}}$ .

**Event  $\mathbf{E}_5$  happens first.** Consider some strategy  $S_5 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_5$  happens. When  $\mathbf{E}_5$  happens, the honest Bob will send  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated. Thus, when  $\mathcal{A}$  sends any transaction containing  $pre_s$ , the  $1 - \alpha$  fraction of honest miners would send  $(pre_s, pre_b, pre_c)$  to  $C_{\text{burn}}^B$ , if they are chosen to mine a block. By Lemma 5.5, if  $C_{\text{burn}}^B$  is activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(-\mathbb{B}c_a^B) + \$E$ . On the other hand, if  $C_{\text{burn}}^B$  is not activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b) + \$E$ . When  $\mathbf{E}_5$  happens, the probability that  $C_{\text{burn}}^B$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, the utility of  $\mathcal{A}$  is at most

$$(1 - \alpha)(\$AV(-\mathbb{B}c_a^B) + \$E) + \alpha(\$AV(\mathbb{B}x_b) + \$E) < \$AV(-\mathbb{A}x_a),$$

where the inequality arises from the fact that  $\$AV(\mathbb{B}c_a^B) > \frac{\$AV(\mathbb{A}x_a + \alpha\mathbb{B}x_b) + \$E}{1 - \alpha}$ .

**Event  $\mathbf{E}_6$  happens first.** Consider some strategy  $S_6 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_6$  happens. When  $\mathbf{E}_6$  happens, the  $1 - \alpha$  fraction of honest miners would send  $pre_a$  to  $A_{\text{burn}}^A$  (or to  $C_{\text{burn}}^A$  if additionally either  $pre_s$  or  $pre_b$  is known), if they are chosen to mine a block. By Lemma 5.5, if either  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) + \$E$ . On the other hand, if neither  $A_{\text{burn}}^A$  nor  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b) + \$E$ . When  $\mathbf{E}_6$  happens, the probability that  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, the utility of  $\mathcal{A}$  is at most

$$(1 - \alpha)(\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) + \$E) + \alpha(\$AV(\mathbb{B}x_b) + \$E) < \$AV(-\mathbb{A}x_a),$$

where the inequality arises from the fact that  $\$AV(\mathbb{A}c_a^A) > \frac{\$AV(\mathbb{B}x_b + \alpha\mathbb{A}x_a) + \$E}{1 - \alpha}$ .

**Event  $\mathbf{E}_7$  happens first.** Consider some strategy  $S_7 \in \bar{\mathcal{R}}$  with non-negligible probability that  $\mathbf{E}_7$  happens. When  $\mathbf{E}_7$  happens, the  $1 - \alpha$  fraction of honest miners would send  $(pre_s, pre_a)$  to  $C_{\text{burn}}^A$ , if they are chosen to mine a block. By Lemma 5.5, if  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) + \$E$ . On the other hand, if  $C_{\text{burn}}^A$  is not activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b) + \$E$ . When  $\mathbf{E}_7$  happens, the probability

that  $C_{\text{burn}}^A$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, by the same calculation as the previous case, the utility of  $\mathcal{A}$  is strictly less than  $\$AV(-\mathbb{A}x_a)$ .

**Event  $E_8$  happens first.** Consider some strategy  $S_8 \in \overline{\mathcal{R}}$  with non-negligible probability that  $E_8$  happens. If Bob enters the execution phase, because Alice does not send  $pre_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ , Bob will send  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated. On the other hand, if Bob enters the abort phase, because Alice does not send  $ping$  to  $P_{\text{refund}}^B$  before AliceChain time  $T_1^A$ , Bob will send  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated. In either case, Bob always sends  $pre_b$  to  $P_{\text{refund}}^B$  as soon as  $A_{\text{defuse}}^A$  is activated.

If  $A_{\text{defuse}}^A$  is guaranteed to be activated at AliceChain time  $t^*$ , it must be activated no later than AliceChain time  $t^* + 1$ . When  $E_8$  happens,  $A_{\text{defuse}}^A$  has been guaranteed to be activated. Thus, when  $E_8$  happens,  $pre_a$  and  $pre_b$  are both publicly known. Then, the  $1 - \alpha$  fraction of honest miners would send  $(pre_a, pre_b)$  to  $C_{\text{burn}}^A$  if they are chosen to mine a block. By Lemma 5.5, if  $C_{\text{burn}}^A$  is activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a - \mathbb{A}c_a^A) + \$E$ . On the other hand, if  $C_{\text{burn}}^A$  is not activated by an honest miner, the utility of  $\mathcal{A}$  is at most  $\$AV(\mathbb{B}x_b) + \$E$ . When  $E_8$  happens, the probability that  $C_{\text{burn}}^A$  is activated by an honest miner is at least  $1 - \alpha$ . Thus, by the same calculation as the previous case, the utility of  $\mathcal{A}$  is strictly less than  $\$AV(-\mathbb{A}x_a)$ .

Finally, notice that the above analysis holds even if  $\mathcal{A}$  may post a new contract on the fly during the protocol execution, since all other players are honest and will not deposit money into the new contract.  $\square$

**Lemma 5.11** (Against Externally Incentivized Alice-Miner Coalition). *Suppose that the hash function  $H(\cdot)$  is a one-way function. Let  $\mathcal{C}$  be a coalition consisting of Bob and possibly any subset of the miners, and let  $\mathcal{A}$  be a disjoint coalition consisting of Alice and at most  $\alpha$  fraction of the mining power where  $\alpha \in [0, 1 - 1/\text{poly}(\lambda)]$ . Suppose that  $\mathcal{C}$  does not have external incentives but  $\mathcal{A}$  may have up to  $\$E$  amount of external incentives. Let  $S_{\mathcal{A}}$  be an arbitrary PPT strategy of  $\mathcal{A}$  that is not in  $\overline{\mathcal{R}}$ . Then, there exists a negligible function  $\text{negl}(\cdot)$  such that except with negligible probability  $\text{negl}(\lambda)$ , it holds that*

$$\text{util}^{\mathcal{C}}(HS_{\mathcal{C}}, S_{\mathcal{A}}, HS_{\mathcal{D}}) \geq 0,$$

where  $\mathcal{D}$  denotes everyone else not in  $\mathcal{C} \cup \mathcal{A}$ .

*Proof.* By Lemma 5.7, any strategy that makes one of  $E_5, E_6, E_7, E_8$  happen is blatantly irrational. By Lemma 5.9, if none of  $E_5, E_6, E_7, E_8$  happen, the one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Refund}^B + \text{Normal}^A)$ , and  $(\text{Refund}^B + \text{Refund}^A)$  must happen. According to Table 3, because  $\$BV(-\mathbb{B}x_b + \mathbb{A}x_a) > 0$ , for all three possible cases, Bob's utility is never negative.  $\square$

#### 5.2.4 Combine Everything Together

Now, we are ready to prove that the atomic swap protocol satisfies bounded maximin fairness.

**Theorem 5.12** (Bounded maximin fairness). *Suppose that  $H(\cdot)$  is a one-way function and suppose the parameters  $\mathbb{B}c_a^B, \mathbb{A}c_a^A, \mathbb{B}c_b^B, \mathbb{A}c_b^A, \alpha \in [0, 1 - 1/\text{poly}(\lambda)]$  satisfy the constraints in Figure 2. Then, the atomic swap protocol satisfies  $\alpha$ -bounded maximin fairness against external incentives.*

*Proof.* Let  $\mathcal{C}$  be a set of honest players where we want to show that  $\mathcal{C}$ 's utility is non-negative, and let  $\mathcal{C}'$  be an externally incentivized coalition. There are five cases to consider.

**Case 1: Alice  $\in \mathcal{C}$  and Bob  $\in \mathcal{C}'$ :** covered by Lemma 5.8.

**Case 2: Bob  $\in \mathcal{C}$  and Alice  $\in \mathcal{C}'$ :** covered by Lemma 5.11.

**Case 3:  $\mathcal{C}$  is miner-only:** It is straightforward to see that no matter how players outside  $\mathcal{C}$  behave, as long as  $\mathcal{C}$  behaves honestly, its utility is non-negative.

**Case 4:  $\text{Alice} \in \mathcal{C}$  and  $\mathcal{C}'$  is miner-only:** Notice that both Alice and Bob are assumed to be honest. In the protocol, Alice and Bob decide whether they will go to the abort phase according to whether Bob sends  $\text{pre}_c$  to  $P_{\text{default}}^B$ . Thus, when Alice and Bob are both honest, both of them enter the execution phase or both of them enter the abort phase. There are two subcases.

- *Subcase 1: Both Alice and Bob enter the execution phase.* Bob only sends  $\text{pre}_c$  to  $P_{\text{default}}^B$  when  $B_{\text{defuse}}^B$  has been activated. Ignoring the negligible probability that  $\mathcal{C}'$  finds  $\text{pre}_c$  by itself,  $B_{\text{burn}}^B$  can never be activated. Alice would send  $\text{pre}_s$  to  $P_{\text{default}}^B$  as soon as she enters the execution phase, and Bob would send  $\text{pre}_s$  to  $P_{\text{default}}^A$  as soon as Alice sent  $\text{pre}_s$  to  $P_{\text{default}}^B$ . Because Alice always sends  $\text{pre}_s$  to  $P_{\text{default}}^B$  before BobChain time  $T_1^B$ , Bob never sends  $\text{pre}_b$ . Besides, Alice never sends  $\text{pre}_a$ . Ignoring the negligible probability that  $\mathcal{C}'$  finds  $\text{pre}_b$  or  $\text{pre}_a$  by itself,  $A_{\text{burn}}^A$ ,  $C_{\text{burn}}^B$  and  $C_{\text{burn}}^A$  can never be activated. When Alice and Bob are in the execution phase,  $P_{\text{refund}}^B$  can be activated only by Bob sending  $\text{pre}_b$ . As we have shown, Bob never sends  $\text{pre}_b$ , so  $P_{\text{refund}}^B$  can never be activated. Moreover, when Alice and Bob are in the execution phase,  $P_{\text{refund}}^A$  can be activated only by Bob sending ping. However, Bob will only send ping to  $P_{\text{refund}}^A$  if  $C_{\text{refund}}^B$  is activated. Because  $P_{\text{refund}}^B$  cannot be activated,  $C_{\text{refund}}^B$  cannot be activated either. Thus,  $P_{\text{refund}}^A$  cannot ever be activated. The  $1 - \alpha$  fraction of honest miner will include Alice's and Bob's transactions, so except the negligible probability,  $P_{\text{default}}^B$  and  $P_{\text{default}}^A$  will be activated in polynomial time. Thus, the honest  $\mathcal{C}$  obtains non-negative utility.
- *Subcase 2: Both Alice and Bob go to the abort phase.* In this case, Bob never sends  $\text{pre}_c$ . In the abort phase Alice always sends ping to  $P_{\text{refund}}^B$  at AliceChain time  $T_0^A$ , so Bob never sends  $\text{pre}_b$ . Similarly, Bob always sends ping to  $P_{\text{refund}}^A$  at AliceChain time  $T_0^A$ , so Alice never sends  $\text{pre}_a$ . Moreover, Alice never sends  $\text{pre}_s$ . Ignoring the negligible probability that  $\mathcal{C}'$  finds  $\text{pre}_s, \text{pre}_a, \text{pre}_b, \text{pre}_c$  by itself,  $A_{\text{burn}}^A$ ,  $B_{\text{burn}}^B$ ,  $P_{\text{default}}^B$ ,  $C_{\text{burn}}^B$ ,  $P_{\text{default}}^A$ ,  $C_{\text{burn}}^A$  can never be activated.

If  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.) has not been active, Alice and Bob send the withdrawal transaction to  $\text{CONTRACT}^B$  ( $\text{CONTRACT}^A$ , resp.). The  $1 - \alpha$  fraction of honest miner will include Alice's and Bob's withdrawal transactions, so except the negligible probability, they can get their deposit back unless the contract becomes active. Next, we analyze different cases depending on which contract becomes active.

- $\text{CONTRACT}^B$  is active. As we have shown,  $B_{\text{burn}}^B$ ,  $P_{\text{default}}^B$ ,  $C_{\text{burn}}^B$  cannot be activated. The  $1 - \alpha$  fraction of honest miner will include Alice's transaction, ping to  $P_{\text{refund}}^B$ , once they mine a block. Thus,  $P_{\text{refund}}^B$  will be activated in polynomial time. When  $\tau^B$  BobChain time has passed since  $P_{\text{refund}}^B$  is activated, Alice and Bob will send ping to  $C_{\text{refund}}^B$ . Again, the  $1 - \alpha$  fraction of honest miner will include Alice's and Bob's transactions, ping to  $C_{\text{refund}}^B$ , once they mine a block. Thus,  $C_{\text{refund}}^B$  will be activated in polynomial time.
- $\text{CONTRACT}^A$  is active. As we have shown,  $A_{\text{burn}}^A$ ,  $P_{\text{default}}^A$ ,  $C_{\text{burn}}^A$  cannot be activated. The  $1 - \alpha$  fraction of honest miner will include Bob's transaction, ping to  $P_{\text{refund}}^A$ , once they mine a block. Thus,  $P_{\text{refund}}^A$  will be activated in polynomial time. When  $\tau^A$  AliceChain time has passed since  $P_{\text{refund}}^A$  is activated, Alice and Bob will send ping to  $C_{\text{refund}}^A$ . Again, the  $1 - \alpha$  fraction of honest miner will include Alice's and Bob's transactions, ping to  $C_{\text{refund}}^A$ , once they mine a block. Thus,  $C_{\text{refund}}^A$  will be activated in polynomial time.

In all cases,  $\mathcal{C}$ 's utility is non-negative except with negligible probability.

**Case 5:  $\mathbf{Bob} \in \mathcal{C}$  and  $\mathcal{C}'$  is miner-only:** The argument is the same as Case 4.  $\square$

### 5.3 Achieving Dropout Resilience

**Theorem 5.13** (Dropout resilience of atomic swap). *Suppose that  $H(\cdot)$  is a one-way function and that all players are PPT machines. Our atomic swap protocol is dropout resilient. In other words, suppose at least  $1/\text{poly}(\lambda)$  fraction of the mining power is honest on either chain; if either Alice or Bob plays honestly but drops out before the end of the protocol, then with  $1 - \text{negl}(\lambda)$  probability, the other party's utility must be non-negative.*

*Proof.* Throughout the proof, for any  $X \in \{pre_s, pre_a, pre_b, pre_c\}$ , we ignore the negligible probability that the miners can find the preimage  $X$  by itself if Alice and Bob have never sent  $X$  before.

We first analyze the cases where Alice drops out. If any of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7, \mathbf{E}_8$  is about to happen, it must be that Alice deviates from the protocol and is about to send a transaction that does not follow from the honest protocol. Because Alice is honest, no matter when Alice drops out, none of  $\mathbf{E}_5, \mathbf{E}_6, \mathbf{E}_7, \mathbf{E}_8$  would happen. By Lemma 5.9, one and only one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Refund}^B + \text{Normal}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$  will happen. According to Table 3, because  $\$BV(\mathbb{A}x_a - \mathbb{B}x_b) > 0$ , for any of the three cases, Bob's utility is non-negative.

Next, we analyze the cases where Bob drops out. If any of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3, \mathbf{E}_4$  is about to happen, it must be that Bob deviates from the protocol and is about to send a transaction that does not follow from the honest protocol. Because Bob is honest, no matter when Bob drops out, none of  $\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3, \mathbf{E}_4$  would happen. By Lemma 5.6, one and only one of  $(\text{Normal}^B + \text{Normal}^A)$ ,  $(\text{Normal}^B + \text{Refund}^A)$  and  $(\text{Refund}^B + \text{Refund}^A)$  will happen. According to Table 4, because  $\$AV(\mathbb{B}x_b - \mathbb{A}x_a) > 0$ , for any of the three cases, Alice's utility is non-negative.  $\square$

## 6 Rational Defection from Grand Coalition

In Theorem 4.1, we prove that the atomic swap protocol is secure assuming at least  $1/\text{poly}(\lambda)$  fraction of honest miners. Here, we justify this assumption by considering a metagame that captures the formation of the coalition. In more detail, we argue that when everyone else joins the coalition, a rational miner will prefer to defect the coalition. In other words, a grand coalition consisting 100% of the mining power is unstable if each individual miner is rational.

### 6.1 Disincentivizing the Grand Coalition Absent of External Incentive.

We first assume the coalition is formed without any external incentive, and we show that any grand coalition whose strategic utility is higher than the honest utility is not stable, and thus justify the assumption of  $1/\text{poly}(\lambda)$  fraction of honest miners for CSP-fairness. The analysis is similar to the metagame for knowledge-coin exchange in [CMST22b].

We start with Alice-miner coalition. According to Lemma 5.2, their utility is better than the honest case only if  $\text{Normal}^B$  and  $\text{Refund}^A$  both happen, which leads to a strategic utility  $\mathbb{B}x_b$ . Suppose Alice controls 100% of the mining power, and they adopt a strategy that leads to  $\text{Normal}^B$  and  $\text{Refund}^A$  with non-negligible probability. When  $\text{Normal}^B$  happens, i.e.,  $P_{\text{default}}^B$  is activated, Bob must have sent  $pre_c$  to  $P_{\text{default}}^B$ , and it implies  $\text{CONTRACT}^A$  has been active and Bob enters the execution phase. If  $\text{CONTRACT}^A$  is active,  $(P_{\text{refund}}^A + C_{\text{refund}}^A)$  is the only possibility for  $\text{Refund}^A$ . However, in the execution phase, Bob sends  $\text{ping}$  to  $P_{\text{default}}^A$  only after  $P_{\text{default}}^B$  has been activated. Thus, the fact that  $\text{Normal}^B$  and  $\text{Refund}^A$  both happen implies that  $P_{\text{refund}}^A$  is activated by  $pre_a$ . When Bob enters the execution phase, there are two possible cases:

- *Alice sends  $pre_s$  to  $P_{\text{default}}^B$  before  $T_1^B$ .* Because  $P_{\text{refund}}^A$  can only be activated after  $T_1^A > T_1^B$ , the moment that  $P_{\text{refund}}^A$  is activated, both  $pre_s$  and  $pre_a$  are publicly known. Within the time window  $\tau^A$  that  $C_{\text{refund}}^A$  can be activated, anyone can send  $(pre_s + pre_a)$  to  $C_{\text{burn}}^A$ .
- *Alice does not send  $pre_s$  before  $T_1^B$ .* In this case, either  $A_{\text{defuse}}^A$  has not been activated, or Bob has sent  $pre_b$  to  $P_{\text{refund}}^B$ . If  $P_{\text{refund}}^A$  is activated, anyone can either send  $pre_a$  to  $A_{\text{burn}}^A$  or send  $(pre_a + pre_b)$  to  $C_{\text{burn}}^A$ .

Recall that the coalition's strategic utility is at most  $\mathbb{B}x_b$ . Suppose the coalition distributes the strategic gain proportionally to the mining power, and consider a small miner  $i$  that has a small mining power  $\gamma$ . If  $i$  joins the coalition and cooperates, its expected reward is at most  $p\gamma \cdot \mathbb{B}x_b$ , where  $p$  denotes the probability that  $P_{\text{refund}}^A$  is activated. On the other hand, suppose  $i$  chooses to not join the coalition. Since its influence to the block generation process is small, we may assume that  $P_{\text{refund}}^A$  is activated with probability  $p$  or more. Without loss of generality, we may assume that every miner in the coalition commits to starving  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  in every block they mine, e.g., by placing a collateral that it will honor its commitment — if not, then the coalition will not be stable since a coalition member will be incentivized to defect from the coalition and claim  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$  itself, which is what we want to prove.

As we have shown above, since the moment  $P_{\text{refund}}^A$  is activated, miner  $i$  has a  $\tau^A$  lead in time to mine a block in which  $i$  can redeem  $\mathbb{B}\epsilon$  from  $A_{\text{burn}}^A$  or  $C_{\text{burn}}^A$ . The probability that  $i$  mines a block in a window of  $\tau^A$  blocks is  $1 - (1 - \gamma)^{\tau^A}$ . Therefore, if  $i$  does not join the coalition, its expected gain would be at least  $p \cdot \mathbb{B}\epsilon \cdot 1 - (1 - \gamma)^{\tau^A}$ . If  $i$  joins the coalition, its expected gain is  $p\gamma \cdot \mathbb{B}x_b$ . Thus, as long as  $p \cdot \mathbb{B}\epsilon \cdot (1 - (1 - \gamma)^{\tau^A}) > p\gamma \cdot \mathbb{B}x_b$ ,  $i$ 's best strategy is to not join the coalition. This means that if everyone else joins the coalition, some small miner  $i$  wants to defect.

Next, we analyze Bob-miner coalition. By Lemma 5.2, the coalition's utility is better off only if  $\text{Refund}^B$  and  $\text{Normal}^A$  both happen.  $\text{Normal}^A$  happens, i.e.,  $P_{\text{default}}^A$  is activated, only if Bob sends  $pre_s$  or Alice sends ping. However, Alice sends ping to  $P_{\text{default}}^A$  only if  $P_{\text{default}}^B$  has been activated which leads to  $\text{Normal}^B$ . Thus, we assume that  $P_{\text{default}}^A$  is activated by  $pre_s$ . Ignoring the negligible probability that Bob finds  $pre_s$  by itself, Alice only sends  $pre_s$  if Bob has sent  $pre_c$  to  $P_{\text{default}}^B$  and  $\text{CONTRACT}^B$  is active. Thus, the fact that  $\text{Refund}^B$  happens implies that  $P_{\text{refund}}^B$  is activated by  $pre_b$ . Consequently, whenever  $P_{\text{refund}}^B$  is activated,  $pre_s$ ,  $pre_b$ , and  $pre_c$  are all publicly known, and anyone can send all of them to  $C_{\text{burn}}^B$ . By the similar argument as Alice-miner coalition, if everyone else joins the coalition, some small miner  $i$  controlling  $\gamma$  fraction of the mining power wants to defect.

## 6.2 Disincentivizing the Grand Coalition in the Presence of External Incentive.

Now we consider the strategic coalition with external incentive, and we show that any grand coalition that harms the honest players is not stable, and thus justify the assumption of  $1/\text{poly}(\lambda)$  fraction of honest miners for bounded maximin fairness.

Let  $\mathbb{E}$  be the upper bound of the external incentive, and let  $\mathbb{V}$  be the maximum value of the atomic swap protocol to the coalition, where  $\mathbb{V} \leq \mathbb{E} + \mathbb{B}x_b$  to Alice-miner coalition and  $\mathbb{V} \leq \mathbb{E} + \mathbb{A}x_a$  to Bob-miner coalition. We start with Alice-miner coalition. According to Table 3, the honest Bob's utility is negative if any of the following conditions hold: 1)  $\text{Refund}^A$  and  $\text{Normal}^B$  both happen; 2)  $\text{Burn}^A$  happens; 3)  $\text{Burn}^B$  happens.

By the choice of  $\mathbb{A}c_a^A$  and  $\mathbb{B}c_a^B$  (Figure 2), once  $\text{Burn}^A$  or  $\text{Burn}^B$  happens, the coalition's utility is negative even if it is compensated by the external incentive. Thus, the only profitable strategy for the externally incentivized coalition that may harm honest players is to invoke  $\text{Refund}^A$  and  $\text{Normal}^B$ .

In this case, we can follow the same argument as in Section 6.1. As long as  $p \cdot \$\epsilon \cdot (1 - (1 - \gamma)^{\tau^A}) > p\gamma \cdot \$V$ , the small miner  $i$ 's best strategy is to not join the coalition.

Similarly, we can show that Bob-miner coalition is not stable. According to Table 4, the honest Alice's utility is negative if any of the following conditions hold: 1)  $\text{Refund}^B$  and  $\text{Normal}^A$  both happen; 2)  $\text{Burn}^A$  happens; 3)  $\text{Burn}^B$  happens. By the choice of  $\mathbb{A}c_b^A$  and  $\mathbb{B}c_b^B$ , once  $\text{Burn}^A$  or  $\text{Burn}^B$  happens, the coalition's utility is negative even if it is compensated by the external incentive. Thus, the only profitable strategy for the externally incentivized coalition that may harm honest players is to invoke  $\text{Refund}^B$  and  $\text{Normal}^A$ . In this case, we can follow the same argument as in Section 6.1, as long as  $p \cdot \$\epsilon \cdot (1 - (1 - \gamma)^{\tau^A}) > p\gamma \cdot \$V$ , the small miner  $i$ 's best strategy is to not join the coalition.

## 7 Instantiation

We now discuss the instantiation of our scheme. We first implement it using a general-purpose smart contract language, and then show an instantiation in Bitcoin's UTXO model.

### 7.1 Ethereum Instantiation

We implemented our contracts in 340 LoC in Solidity, Ethereum's smart contract language. Transaction fees on Ethereum are determined by gas usage, which corresponds to the total cost of operations executed by the contract.

We give the gas cost of our scheme in Table 5. Additionally, we compare our gas cost to those of Rapidash's CSP-fair swap [CMST22b] in Table 6. As expected, our costs are slightly higher overall due to the extra input logic and the extra defuse logic that are necessary to achieve bounded maximin fairness. However, we note that once the contracts are active, the optimistic paths have similar costs in both schemes.

Table 5: Bounded maximin fair atomic swap, gas cost. (O) denotes an optimistic case where both Alice and Bob are honest.

Contract	Redeem path	Gas
CONTRACT <sup>B</sup>	Input, Alice	48,355
	Input, Bob	50,583
	Withdraw, Alice	35,956
	Withdraw, Bob	38,271
	(O) $(P_{\text{defuse}}^B)$ , Alice	35,405
	(O) $(B_{\text{defuse}}^B + P_{\text{defuse}}^B)$ , Bob	88,399
	Refund $(P_{\text{refund}}^B + C_{\text{refund}}^B)$ , Alice	114,662
	Refund $(B_{\text{defuse}}^B + P_{\text{refund}}^B + C_{\text{refund}}^B)$ , Bob	147,567
	Early Bomb $(B_{\text{burn}}^B)$ , Miner	50,295
	Bomb $(C_{\text{burn}}^B)$ , Miner	57,152
CONTRACT <sup>A</sup>	Input, Alice	50,650
	Input, Bob	48,333
	Withdraw, Alice	38,251
	Withdraw, Bob	35,912
	(O) $(P_{\text{defuse}}^A)$ , Alice	54,925
	(O) $(P_{\text{defuse}}^A)$ , Bob	58,679
	Refund $(A_{\text{defuse}}^A + P_{\text{refund}}^A + C_{\text{refund}}^A)$ , Alice	149,634
	Refund $(A_{\text{defuse}}^A + P_{\text{refund}}^A + C_{\text{refund}}^A)$ , Bob	145,894
	Early bomb $(A_{\text{burn}}^A)$ , Miner	49,956
	Bomb $(C_{\text{burn}}^A)$ , Miner	53,475

Table 6: Gas cost comparison between Rapidash and our bounded maximin fair atomic swap. (O) denotes an optimistic case.

Contract	Redeem path	Gas
RAPIDASH <sup>B</sup>	Normal path ( $P_{\text{default}}^B$ ), Alice	52,279
	Normal path ( $P_{\text{default}}^B$ ), Bob	56,681
	Refund path ( $P_{\text{refund}}^B + C_{\text{refund}}^B$ ), Bob	123,631
	Burn path ( $C_{\text{burn}}^B$ ), Miner	42,266
RAPIDASH <sup>A</sup>	Input, Alice	50,465
	Input, Bob	55,817
	Withdraw, Alice	38,228
	Withdraw, Bob	35,911
	(O) ( $P_{\text{default}}^A$ ), Alice	54,904
	(O) ( $P_{\text{default}}^A$ ), Bob	58,656
	Refund ( $P_{\text{refund}}^A + C_{\text{refund}}^A$ ), Alice	118,379
	Refund ( $P_{\text{refund}}^A + C_{\text{refund}}^A$ ), Bob	114,647
	Burn ( $C_{\text{burn}}^A$ ), Miner	53,431
	Input, Alice	48,355
CONTRACT <sup>B</sup>	Input, Bob	50,583
	Withdraw, Alice	35,956
	Withdraw, Bob	38,271
	(O) ( $P_{\text{default}}^B$ ), Alice	35,405
	(O) ( $B_{\text{defuse}}^B + P_{\text{default}}^B$ ), Bob	88,399
	Refund ( $P_{\text{refund}}^B + C_{\text{refund}}^B$ ), Alice	114,662
	Refund ( $B_{\text{defuse}}^B + P_{\text{refund}}^B + C_{\text{refund}}^B$ ), Bob	147,567
	Early Bomb ( $B_{\text{burn}}^B$ ), Miner	50,295
	Bomb ( $C_{\text{burn}}^B$ ), Miner	57,152
	Input, Alice	50,650
CONTRACT <sup>A</sup>	Input, Bob	48,333
	Withdraw, Alice	38,251
	Withdraw, Bob	35,912
	(O) ( $P_{\text{default}}^A$ ), Alice	54,925
	(O) ( $P_{\text{default}}^A$ ), Bob	58,679
	Refund ( $A_{\text{defuse}}^A + P_{\text{refund}}^A + C_{\text{refund}}^A$ ), Alice	149,634
	Refund ( $A_{\text{defuse}}^A + P_{\text{refund}}^A + C_{\text{refund}}^A$ ), Bob	145,894
	Early bomb ( $A_{\text{burn}}^A$ ), Miner	49,956
	Bomb ( $C_{\text{burn}}^A$ ), Miner	53,475

## 7.2 Bitcoin Instantiation

Traditional smart contracts operate by receiving coins from users, holding them within the contract until specific conditions activate the transfer. Once triggered, the contract disperses some or all of the coins to designated recipients. Bitcoin, however, follows the Unspent Transaction Output (UTXO) model, where coins are linked to addresses represented as  $Adr \in \{0, 1\}^\lambda$ . Each address can be spent only once, ensuring that once used in a transaction, it cannot be reused. Transactions are recorded on the blockchain, facilitating the movement of coins from input addresses to new output addresses. Any excess coins not explicitly allocated are collected by the miner as a transaction fee for processing the block.

Bitcoin transactions can be thought of as being generated by the transaction function  $tx$ . A transaction  $tx_A$ , denoted

$$tx_A := tx \left( \begin{array}{l} [(Adr_1, \Phi_1, \$v_1), \dots, (Adr_n, \Phi_n, \$v_n)], \\ [(Adr'_1, \Phi'_1, \$v'_1), \dots, (Adr'_m, \Phi'_m, \$v'_m)] \end{array} \right),$$

transfers  $v_i$  coins from each input address  $Adr_i$ ,  $i \in [n]$ , and deposits  $v'_i$  coins to each output address  $Adr'_j$ ,  $j \in [m]$ . To prevent minting coins from thin air, it must be guaranteed that  $\sum_{i \in [n]} \$v_i \geq \sum_{j \in [m]} \$v'_j$ , where the difference is offered as transaction fee to the miners.

Bitcoin addresses are governed by *scripts*  $\Phi : \{0, 1\}^\lambda \rightarrow \{0, 1\}$ , which define spending conditions for their coins. Unlike smart contracts, which allow arbitrary logic, Bitcoin's scripting language has limited expressiveness. A transaction is *authorized* when witnesses  $[x_1, \dots, x_n]$  satisfy  $\Phi_i(x_i) = 1$  for all  $i$ , and *confirmed* once recorded on the blockchain.

Since addresses can only be spent *once*, contract logic must be encoded directly within scripts. Our approach relies on standard Bitcoin scripts, ensuring compatibility. An address  $Adr$  is controlled by a user if they possess a valid witness  $x$  such that  $\Phi(x) = 1$ .

### 7.2.1 Instantiating Contract<sup>B</sup> with Bounded Maximin Fairness

We depict the payment flow in Figure 7. Alice and Bob use transaction  $tx_{\text{stp}}^B$  that creates  $Adr_{\text{stp}}^B$  with  $\$x_b + \$c_a^B + \$c_b^B$  coins and another address  $Adr_{B_{\text{defuse}}^B}$  with  $\$n^B$  coins in it. Here  $\$x_b + \$c_b^B + \$n^B$  coins come from Bob and  $\$c_a^B$  coins come from Alice. We have two new transactions  $tx_{B_{\text{defuse}}^B}$  and  $tx_{B_{\text{burn}}^B}$ , corresponding to the activation points  $B_{\text{defuse}}^B$  and  $B_{\text{burn}}^B$ , respectively. Transaction  $tx_{B_{\text{defuse}}^B}$  redeems  $\$n^B$  coins from the address  $Adr_{B_{\text{defuse}}^B}$  provided a timeout of  $T_1^B$  has passed since the setup transaction was published on the blockchain. The other transaction  $tx_{B_{\text{burn}}^B}$  redeems the  $\$n^B$  coins from  $Adr_{B_{\text{defuse}}^B}$ , as well as the coins from the setup address  $Adr_{\text{stp}}^B$  provided  $pre_c$  is released. The transaction burns  $(\$x_b + \$c_a^B + \$c_b^B + \$n^B - \$\epsilon^B)$  coins leaving behind  $\$e^B$  coins as transaction fees.

Transaction  $tx_{P_{\text{default}}^B}$  can only be activated with  $pre_s$  (pre-image of  $h_s$ ) and  $z_c$  (pre-image of  $h_c$ ) are available. This transaction spends coins from  $Adr_{\text{stp}}^B$ , and sends  $\$x_b + \$c_a^B$  coins to an address of Alice and send  $\$c_b^B$  coins to an address of Bob. We then have two cases: two transactions  $tx_{P_{\text{refund}}^B}$  and  $tx_{P_{\text{refund}}^B}^{\text{ping}}$  can activate  $P_{\text{refund}}^B$  and spend from  $Adr_{\text{stp}}^B$  after a timeout of  $T_1^B$  has passed. The only difference is  $tx_{P_{\text{refund}}^B}$  can be activated by Bob only with pre-image  $pre_b$ , while anyone, especially Alice can publish  $tx_{P_{\text{refund}}^B}^{\text{ping}}$  without needing any pre-image. If either one of those transactions are published, after a timeout of  $\tau^B$ , Alice and Bob can get a full refund of their locked coins with transactions  $tx_{C_{\text{refund}}^B}$  and  $tx_{C_{\text{refund}}^B}^{\text{ping}}$ , respectively. However, before the timeout, the  $C_{\text{burn}}^B$  path can be activated in both cases before the timeout  $\tau^B$  passes using  $tx_{C_{\text{burn}}^B}^{P_{\text{refund}}^B}$  and  $tx_{C_{\text{burn}}^B}^{P_{\text{refund}}^B, \text{ping}}$ , respectively. In

both cases, all of  $pre_s, pre_b, pre_c$  must be available. Finally we have transaction  $tx_{C_{burn}^B}$  that can also directly activate  $C_{burn}^B$  path from  $Adr_{stp}^B$  given all of  $pre_s, pre_b, pre_c$ . All transactions must be signed by both Alice and Bob.

**Protocol Flow.** Alice and bob first agree on the setup transaction  $tx_{stp}^B$  and sign all other redeeming transactions. They ensure that the signatures on the transaction  $tx_{B_{defuse}^B}$  are only with Bob, and the signatures on  $tx_{P_{refund}^B}^{ping}$  are only with Alice. Alice and Bob keep these signatures privately and broadcast all other transactions and signatures into the network. By posting the setup transaction on the blockchain, we enter the execution phase.

Whenever Bob wishes to activate  $B_{defuse}^B$  branch, he publishes  $tx_{B_{defuse}^B}$  along with the corresponding signatures on the blockchain. Notice that since  $tx_{B_{defuse}^B}$  and  $tx_{B_{burn}^B}$  spend from the address  $Adr_{B_{defuse}^B}$ , they are mutually exclusive, meaning only one of them can be posted on the blockchain. Notice that any of the other branches can be activated if and only if  $B_{defuse}^B$  was activated (or in other words,  $B_{burn}^B$  was not activated with  $tx_{B_{burn}^B}$ ). Alice can publish  $tx_{P_{default}^B}$  with pre-images  $pre_s$  and  $pre_c$ . If not, Bob can initiate the refund using  $tx_{P_{refund}^B}$  after a timeout of  $T_1^B$  and using the pre-image  $pre_b$ . Whenever Alice wants to activate  $P_{refund}^B$  with a ping, she publishes the transaction  $tx_{P_{refund}^B}^{ping}$  along with valid signatures in her possession. If the transaction is published, activation point  $C_{refund}^B$  can be activated by  $tx_{C_{refund}^B}^{ping}$  after a timeout of  $\tau^B$  time. The  $C_{burn}^B$  path can be activated using transactions  $tx_{C_{burn}^B}$ ,  $tx_{C_{burn}^B}^{P_{refund}^B}$  or  $tx_{C_{burn}^B}^{P_{refund}^B, ping}$  depending on which paths have been activated so far.

### 7.2.2 Instantiating Contract<sup>A</sup> with Bounded Maximin Fairness

The payment flow is depicted in Figure 8. Alice and Bob use transaction  $tx_{stp}^A$  that creates  $Adr_{stp}^A$  with  $\$x_a + \$c_a^A + \$c_b^A$  coins and another address  $Adr_{A_{defuse}^A}$  with  $\$y^A$  coins in it. Here  $x_a + \$c_a^A + \$y^A$  coins come from Alice and  $\$c_b^A$  coins come from Bob. We have two new transactions  $tx_{A_{defuse}^A}$  and  $tx_{A_{burn}^A}$ , corresponding to the activation points  $A_{defuse}^A$  and  $A_{burn}^A$ , respectively. Transaction  $tx_{A_{defuse}^A}$  redeems  $\$y^B$  coins from the address  $Adr_{A_{defuse}^A}$  provided a timeout of  $T_1^A$  has passed since the setup transaction was published on the blockchain. The other transaction  $tx_{A_{burn}^A}$  redeems the  $\$y^A$  coins from  $Adr_{A_{defuse}^A}$ , as well as the coins from the setup address  $Adr_{stp}^A$  provided  $pre_a$  or  $pre_b$  is released. The transaction burns  $(\$x_a + \$c_a^A + \$c_b^A + \$y^A - \$\epsilon^A)$  coins leaving behind  $\$y^A$  coins as transaction fees.

We have transaction  $tx_{P_{default}^A}^{ping}$  that can activate  $P_{default}^A$  by a ping by anyone. On the other hand, we have transaction  $tx_{P_{default}^A}$  that can only be activated with  $pre_s$  (pre-image of  $h_s$ ) is available. Both of these transactions will spend the coins from  $Adr_{stp}^A$ , and send  $\$x_a + \$c_a^A$  coins to an address of Bob and send  $\$c_b^A$  coins to an address of Alice. We then have two cases: two transactions  $tx_{P_{refund}^A}$  and  $tx_{P_{refund}^A}^{ping}$  can activate  $P_{refund}^A$  and spend from  $Adr_{stp}^A$  after a timeout of  $T_1^A$  has passed. The only difference is  $tx_{P_{refund}^A}$  can be activated by Alice only with pre-image  $pre_a$ , while anyone, especially Bob can publish  $tx_{P_{refund}^A}^{ping}$  without needing any pre-image. If either one of those transactions are published, after a timeout of  $\tau^B$ , Alice and Bob can get a full refund of their locked coins with transactions  $tx_{C_{refund}^A}$  and  $tx_{C_{refund}^A}^{ping}$ , respectively. However, before the timeout, the  $C_{burn}^A$  path can be activated in both cases using  $tx_{C_{burn}^A}^{P_{refund}^A}$  and  $tx_{C_{burn}^A}^{P_{refund}^A, ping}$ , respectively. In both cases, either  $pre_a$  and

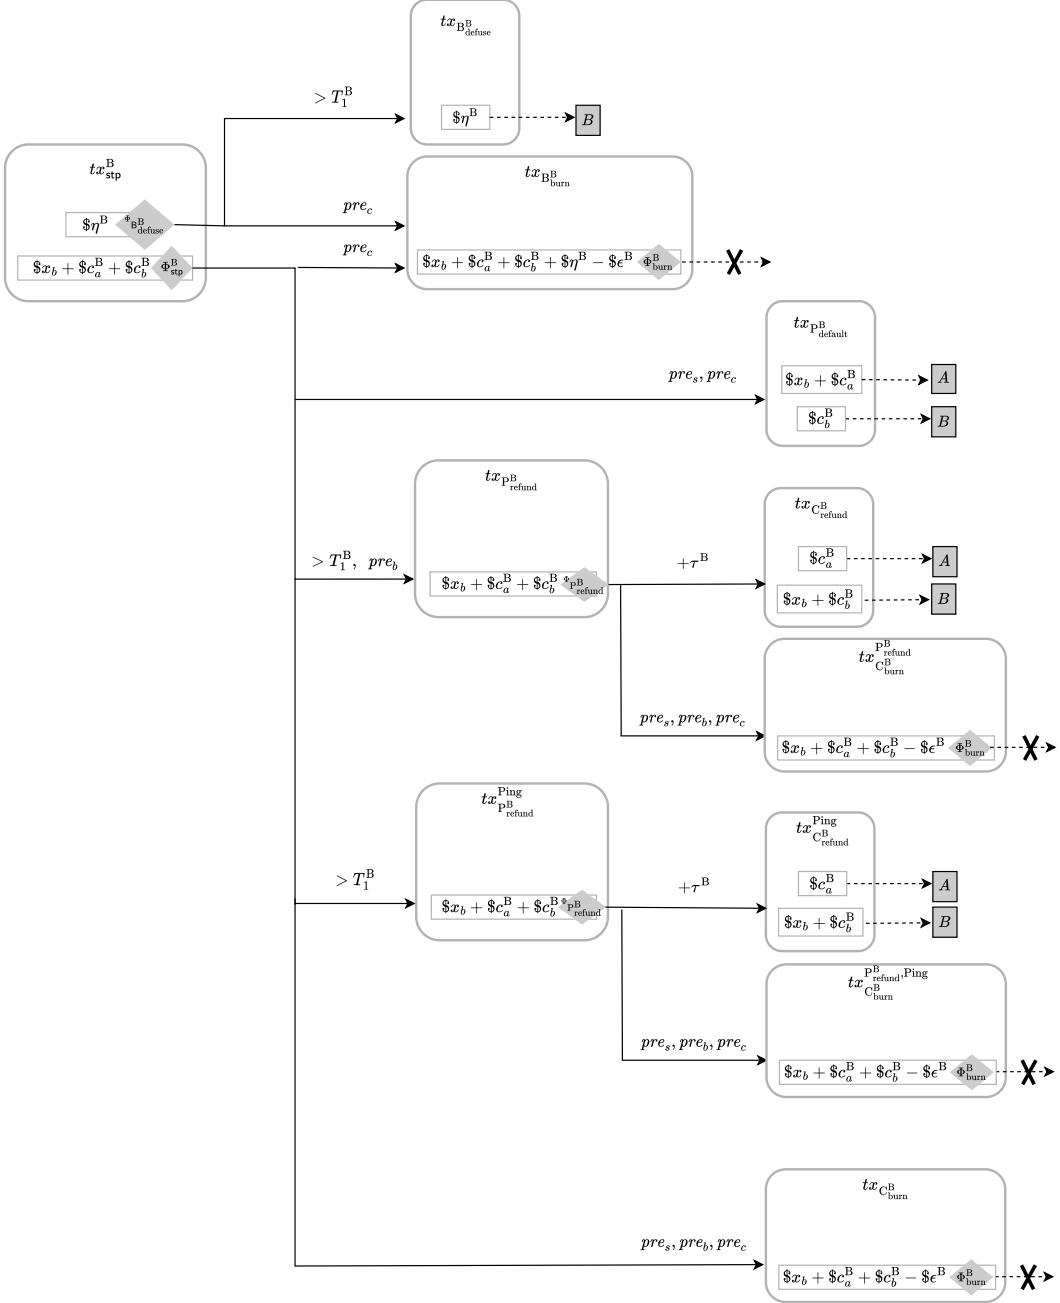


Figure 7: The transaction flow of  $\text{CONTRACT}^B$  in Bitcoin for an atomic swap with bounded maximin fairness. Rounded boxes denote transactions, and rectangles within are outputs of the transaction. Incoming arrows denote transaction inputs, outgoing arrows denote how an output can be spent by a transaction at the end of the arrow. Solid lines indicate the transaction output can be spent only if both users sign the spending transaction. Dashed arrows indicate that the output can be spent by one user ( $A$  for Alice and  $B$  for Bob).

$pre_s$ , or  $pre_a$  and  $pre_b$  must be available. Finally we have transaction  $tx_{C_{burn}^A}$  that can also directly activate  $C_{burn}^A$  path from  $Adr_{stp}^A$  given either  $pre_a$  and  $pre_s$ , or  $pre_a$  and  $pre_b$  are available. All

transactions are required to be signed by both Alice and Bob.

**Protocol Flow.** We proceed exactly as in  $\text{CONTRACT}^B$ , except for the two additional transactions in  $tx_{A_{\text{defuse}}^A}$  and  $tx_{A_{\text{burn}}^A}$ . Alice and Bob sign the two transactions prior to signing the setup transaction. Alice has the signatures on  $tx_{A_{\text{defuse}}^A}$  that she keeps privately, while  $tx_{A_{\text{burn}}^A}$  and the signatures on this transaction are broadcast to the network. After the setup transaction,  $tx_{\text{stp}}^A$  is published on the blockchain, and the execution phase begins whenever  $A_{\text{defuse}}^A$  is to be activated, Alice publishes  $tx_{A_{\text{defuse}}^A}$  and the signatures on the blockchain. To activate  $A_{\text{burn}}^A$ , transaction  $tx_{A_{\text{burn}}^A}$  along with the corresponding signatures, and either  $pre_a$  or  $pre_s$  are published on the blockchain. Notice that, as required, only one of these two transactions can be posted, allowing us to realize that  $A_{\text{defuse}}^A$  and  $A_{\text{burn}}^A$  are mutually exclusive. The rest of the protocol proceeds as the description for the previous instantiation.

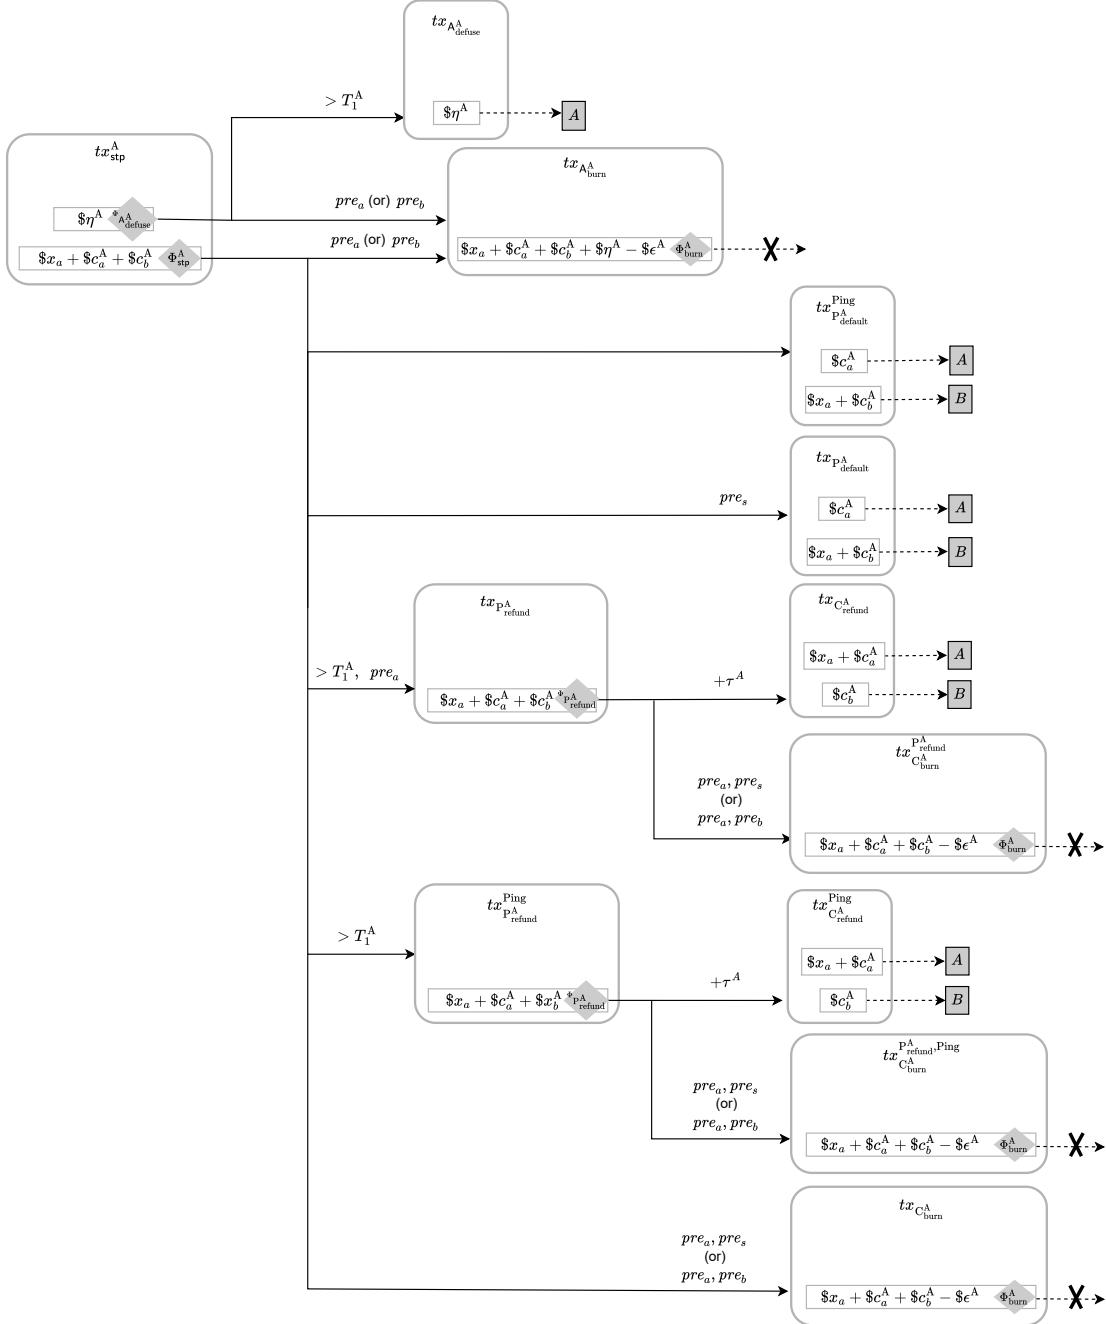


Figure 8: The transaction flow of  $\text{CONTRACT}^A$  in Bitcoin for an atomic swap with bounded maximin fairness. Rounded boxes denote transactions, and rectangles within are outputs of the transaction. Incoming arrows denote transaction inputs, outgoing arrows denote how an output can be spent by a transaction at the end of the arrow. Solid lines indicate the transaction output can be spent only if both users sign the spending transaction. Dashed arrows indicate that the output can be spent by one user ( $A$  for Alice, and  $B$  for Bob). The timelock ( $T_1^A$  and  $\tau^A$ ) associated with a transaction output is written over the corresponding outgoing arrow.

## References

- [Bon16] Joseph Bonneau. Why buy when you can rent? - bribery attacks on bitcoin-style consensus. In *FC Workshops*, 2016.
- [CGJ<sup>+</sup>17] Arka Rai Choudhuri, Matthew Green, Abhishek Jain, Gabriel Kaptchuk, and Ian Miers. Fairness in an unfair world: Fair multiparty computation from public bulletin boards. In *ACM CCS*, 2017.
- [CGL<sup>+</sup>18] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, volume 11239, pages 563–596, 2018.
- [CMST22a] Hao Chung, Elisaweta Masserova, Elaine Shi, and Sri AravindaKrishnan Thyagarajan. Ponyta: Foundations of side-contract-resilient fair exchange. *Cryptology ePrint Archive*, 2022.
- [CMST22b] Hao Chung, Elisaweta Masserova, Elaine Shi, and Sri AravindaKrishnan Thyagarajan. Rapidash: Foundations of side-contract-resilient fair exchange. *Cryptology ePrint Archive*, 2022.
- [CRS24] Hao Chung, Tim Roughgarden, and Elaine Shi. Collusion-resilience in transaction fee mechanism design. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1045–1073, 2024.
- [CS23] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *SODA*, 2023.
- [DH20] Apoorvaa Deshpande and Maurice Herlihy. Privacy-preserving cross-chain atomic swaps. In *International conference on financial cryptography and data security*, pages 540–549. Springer, 2020.
- [Her18] Maurice Herlihy. Atomic cross-chain swaps. In *PODC*, 2018.
- [MD19] Mahdi H. Miraz and David C. Donald. Atomic cross-chain swaps: Development, trajectory and potential of non-monetary digital token swap facilities. In *AETiC*, 2019.
- [MHM18] Patrick McCorry, Alexander Hicks, and Sarah Meiklejohn. Smart contracts for bribing miners. In *FC Workshops*, 2018.
- [MMS<sup>+</sup>] Giulio Malavolta, Pedro Moreno-Sánchez, Clara Schneidewind, Aniket Kate, and Matteo Maffei. Anonymous multi-hop locks for blockchain scalability and interoperability. In *NDSS 2019*.
- [ora22] Oracle manipulation attacks are rising, creating a unique concern for defi. 2022.
- [PG99] Henning Pagnia and Felix C. Gartner. On the impossibility of fair exchange without a trusted third party. Technical report, 1999.
- [PS17] Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *PODC*, 2017.
- [SCW23] Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design? In *ITCS 2023*, 2023.

[TSW24] Sri Aravinda Krishnan Thyagarajan, Pratik Soni, and Ke Wu. Game-theoretically fair distributed sampling. In *CRYPTO (8)*, volume 14927 of *Lecture Notes in Computer Science*, pages 207–239. Springer, 2024.

[TYME21] Itay Tsabary, Matan Yechieli, Alex Manuskin, and Ittay Eyal. MAD-HTLC: because HTLC is crazy-cheap to attack. In *S&P*, 2021.

[vdM19] Ron van der Meyden. On the specification and verification of atomic swap smart contracts. In *IEEE ICBC*, 2019.

[WAS22] Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. In *Eurocrypt*, 2022.

[WHF19] Fredrik Winzer, Benjamin Herd, and Sebastian Faust. Temporary censorship attacks in the presence of rational miners. In *IEEE European Symposium on Security and Privacy Workshops*, 2019.

[WSZN22] Sarisht Wadhwa, Jannis Stoeter, Fan Zhang, and Kartik Nayak. He-htlc: Revisiting incentives in htlc. Cryptology ePrint Archive, Paper 2022/546, 2022. <https://eprint.iacr.org/2022/546>.

[ZDBN19] Jean-Yves Zie, Jean-Christophe Deneuville, Jeremy Briffaut, and Benjamin Nguyen. Extending atomic cross-chain swaps. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, 2019.