

RESEARCH ARTICLE

Transaction Fees, Block Size Limit, and Auctions in Bitcoin

Nicola Dimitri^{*†}

Abstract. Confirmation of Bitcoin transactions is executed in blocks, which are then stored in the Blockchain. As compared to the number of transactions in the *mempool*, the set of transactions which are verified but not yet confirmed, available space for inclusion in a block is typically limited. For this reason, successful miners can only process a subset of such transactions, and users compete with each other to enter the next block by offering confirmation fees. Assuming that successful miners pursue revenue maximization, they will include in the block those *mempool* transactions that maximize earnings from related fees. In the paper we model transaction fees as a Nash Equilibrium outcome of an auction game with complete information. In the game the successful miner acts *as* an auctioneer selling block space, and users bid for shares of such space to confirm their transactions. Moreover, based on expected fees we also discuss what the optimal, revenue maximizing, block size limit should be for the successful miner. Consistently with the intuition, the optimal block size limit resolves the trade-off between including additional transactions (which possibly lower the unit fees collected) and keeping the block capacity limited (with, however, higher unit fees).

1. Introduction

Cryptocurrencies, notably Bitcoin, have recently drawn much attention within the academic community, among investors and the public, for more than one reason, and in particular for the highly variable behavior of the Dollar/Bitcoin exchange rate.

Those cryptocurrencies that follow the original Bitcoin protocol exhibit some differences but also a number of similarities. In particular, all of them share a system of transactions based on two main competitive activities.

The first is among miners, nodes in the network who attempt to solve the outstanding cryptopuzzle. Indeed, its solution would allow miners to gain the reward defined by the protocol and the transaction fees paid by the users, who wish to have their exchanges confirmed as quickly as possible.^{1, 2} While the former type of compensation is known to the miners with certainty, the latter may vary from one block to the next, depending upon the users' willingness to pay for confirming their transactions.³

The second competitive activity takes place among users, aiming to confirm a transaction the soonest. They offer fees to the successful miner for their transaction to be confirmed in the next block. Given the block size limit, the miner will confirm those Bitcoin transfers whose total fees maximize his revenue. If the block could accommodate all transactions pending in the *mempool*—that is, those already verified but not yet confirmed—then revenue maximization would coincide with their full inclusion. However, typically each block could not contain all

* 1N9ukmAq6EhhVigrAHiMMzSHdEwDAcLskP

† Nicola Dimitri (dimitri@unisi.it) is Professor of Economics at the University of Siena, Italy, and Research Associate at the Centre for Blockchain Technologies, University College London.

outstanding transactions, and this is why users compete with each other for space, by proposing fees to confirm their exchanges in the next block.⁴⁻¹⁰

This paper discusses the latter kind of competition, interpreting transaction fees as price bids submitted by the users to obtain (*buy*) a share of the available block space for transaction confirmation. Under such an interpretation, competition for block space *can be seen as an auction*, with the successful miner acting as an *auctioneer*, *selling* space for confirmation, and the proposed transaction fees acting as price bids to obtain *shares* of the available space.¹¹ We are not aware of such an approach to investigate the block size limit and transaction fees, with the closest in the existing literature being perhaps the work by Lavi (*et al.*) (“Redesigning Bitcoin’s Fee Market,” 2017).⁹

Given the block capacity, the transaction size and associated proposed fees will determine which exchanges are confirmed in the next block. For this reason, users compete to enclose their transactions among those maximizing the miner’s revenue. In so doing they face a fundamental trade-off: the higher the fee offered, the more likely it is for a transaction to be confirmed soon, but the more expensive—in case of confirmation—the payment will be. For this reason, the proposed fee is likely to be the outcome of a strategic decision. Namely, how large a fee to offer may depend also on what the other users are expected, or known, to offer.

The problem bears similarities with the well-known Rucksack Problem,¹² though with a main difference. More specifically, in our approach a variation in the block capacity would typically affect the proposed transaction fees, while in the classical Rucksack problem this does not take place. Though outside the scope of this paper, in analogy with the original Rucksack problem it could be interesting to investigate whether the block composition is also an NP-hard problem.

The block size limit is decided by the Bitcoin community. For this reason, in the paper we shall consider the block size limit as given, nevertheless discussing conditions for it to maximize a miner’s revenue. Therefore, we shall model which transaction fees to offer as decisions in a static auction game with complete information, played only by the users and not by the miners. Moreover, in order to focus on the main issue, in this paper we shall not study the effects of possible non-competitive behavior among miners and users.

The findings suggest some interesting insights. The first, main, message of the work is that the revenue maximizing block size limit for the miners depends on the distribution of the users’ willingness to pay for the transaction confirmation, which we call the *value* of a transaction. More specifically, if the values’ distribution is polarized, with some very high and some rather low values, the optimal size limit of a block would be *small*. Alternatively, if the values are uniformly distributed, some *intermediate* size limit would be optimal. Meanwhile, if the values are clustered, and *close* to each other, it would be preferable to have a *large* block size limit.

If it is reasonable to think that the *importance* of a transaction may depend on the underlying amount of paid bitcoins, as well as on the exchange rate of fiat currency to Bitcoin, for example Dollar/Bitcoin. An increase in such rate is likely to increase the proposed transaction fees. This is consistent with recent empirical evidence, where an increase in the Dollar/Bitcoin exchange rate induced a meaningful rise in the fees. However, as we shall see, to focus on our main issue, with no major loss of generality, the exchange rate in the analysis will be kept fixed and therefore irrelevant to our findings. The total number of exchanges can also affect the level of transaction fees, since demand for space by a higher number of users can increase bid competition.

The paper is structured as follows: Section 2 introduces some foundational notions and concepts of the model; Section 3 characterizes transaction fees as a Nash Equilibrium of an auction game for confirmation; and Section 4 presents some further discussion and concludes the paper.

2. The Model Primitives

2.1 Transactions—At each date $t = 0, 1, 2, \dots$, the fundamental unit of analysis is the outstanding generic transaction τ_t in the *mempool*, hence verified but not yet confirmed. A transaction is described by the following vector of elements

$$\tau_t = (x, y, b_\tau, s_\tau, t_\tau \leq t, f_\tau \leq v_\tau, \delta_\tau)$$

where x is the payer and y the payee, b_τ the number of bitcoins exchanged in the transaction, s_τ the transaction size measured in bytes, t_τ the date at which the transaction has been executed, f_τ the transaction fee proposed for its confirmation. Moreover, v_τ is the *value* of the transaction fee, namely the maximum number of bitcoins that the parties in the transaction are willing to pay to the miner, while $0 \leq \delta_\tau \leq 1$ the discount rate associated to the transaction, which could be defined by the lowest discount rate of the two parties in the exchange. The discount rate quantifies the degree of impatience of the parties behind the transaction. If $U_t(\tau_{t_\tau})$ is the utility at time t , for simplicity, of both parties provided by transaction τ executed at t_τ , then

$$U_t(\tau_{t_\tau}) = \delta_\tau^{t-t_\tau} U_{t_\tau}(\tau_{t_\tau}) \quad (1)$$

that is the utility of a transaction decreases exponentially with the time interval $t - t_\tau$ between its execution and the current date. The proposed fee f_τ is determined strategically by the parties and we assume it cannot be larger than v_τ .

In general, if $0 \leq \theta_\tau \leq 1$ is a share of the exchanged sum b_τ , we would expect

$$f_\tau = \theta_\tau b_\tau \leq v_\tau \leq b_\tau$$

and therefore $\theta_\tau \leq \frac{v_\tau}{b_\tau}$. Indeed, in general it is reasonable to think that the fee is smaller than the paid sum of bitcoins. Moreover, the subscript τ of θ_τ indicates that such coefficient may vary across different transactions. However, two observations are worth mentioning. First, that in principle we could not exclude that $\theta_\tau > 1$ and, second, that the coefficient θ_τ in turn may depend upon a number of quantities. In general, we may envisage θ_τ to be defined as the following function θ

$$\theta_\tau = \theta(S_i, n_i, \Delta e_i, b_\tau, \delta_\tau)$$

where S_i stands for the size limit of the next i th block, n_i indicates the number of outstanding transactions and Δe_i the expected variation of the exchange rate $\frac{\text{dollar}}{\text{bitcoin}}$, when confirmation is expected to take place.

Broadly speaking we find it intuitive to assume that, keeping all the rest as given, θ_τ would decrease with S_i and δ_τ and increase with n_i and b_τ . Currently in Bitcoin the block size limit $S_i = S$ is constant over time, commonly known among miners and users though, in principle, S_i could even be a random variable for individuals when offering their fees. That is, in a more general scenario the community may leave completely free the successful miner to announce the block size limit, only after having mined it. Likewise, n_i and Δe_i are also random variables for x when proposing f_τ .

2.2 The Miners' Goal—Although the block size limit (BSL) is decided by the Bitcoin community, it is nonetheless interesting to investigate under what conditions the agreed-upon size limit is optimal for the successful miner. To discuss optimality in this paragraph we first define the miner's goal, in a very general way, that is with BSL being the choice variable. This will allow us to conceptualize the optimal BSL for the miner, and its connection with the transaction fee analysis.

A main intuition behind the discussion on the optimal BSL for miner is the following. The larger BSL, the higher the number of transactions to be included—but, presumably, the lower the proposed transaction fees for each transaction. That is, in general we envisage a trade-off between the fee collected for each confirmed transaction and the block capacity.

Broadly speaking, the miner's problem could be formulated as follows.

Let S be BSL as measured in bytes, and $r(S)$ the expected reward (fee) per byte of confirmed transactions, when BSL is S .

Then the miner's expected revenue $R(S)$, expressed as a function of BSL, is given by $R(S) = Sr(S)$ and the optimal BSL would solve

$$\text{Max}_S R(S) = Sr(S) \text{ with } S > 0 \quad (2)$$

Treating, for simplicity, S as continuous and assuming $r(S)$ to be twice differentiable it follows that

$$\frac{dR(S)}{dS} = r(S) + Sr'(S) \quad (3)$$

while

$$\frac{d^2R(S)}{dS^2} = 2r'(S) + Sr''(S) \quad (4)$$

The solution to (2) would clearly depend on the shape of $r(S)$.

If $r'(S) > 0$ for all S , then (3) is positive and (2) is solved by the largest possible S , that is a BSL which will include all the outstanding transactions. In this case, adding an extra transaction to the block would more than compensate for a possible decrease in the single transaction fee.

If instead the solution to (2) is a stationary point of $R(S)$, then (2) is solved by the first order condition, that is by $S = S^*$ such that

$$r(S^*) + S^*r'(S^*) = 0 \quad (5)$$

hence by

$$\frac{r'(S^*)S^*}{r(S^*)} = -1 \quad (6)$$

At the optimal S^* the function $r(S)$ has *unit elasticity* (in absolute value), that is where the % change of S is exactly counterbalanced by the % change of $r(S)$, though in the opposite direction.

As an illustration, suppose $r(S) = a - bS$ with $a, b > 0$: then the optimal BSL would be $S^* = \frac{a}{2b}$ suggesting that the higher(smaller) $a(b)$ the larger the optimal BSL.

Example 1—Suppose S is the BSL agreed upon within the community, and that there are three outstanding transactions $\tau = 1, 2, 3$, respectively, with the following features ($f_1 = 3$; $s_1 = \frac{2S}{3}$), ($f_2 = 2$; $s_2 = \frac{S}{2}$); ($f_3 = 1,5$; $s_3 = \frac{S}{2}$). In this case revenue maximization requires confirming transactions $\tau = 2$ and $\tau = 3$, since $f_2 + f_3 = 2 + 1,5 = 3,5$ against $f_1 = 3$. Yet, the average (per occupied byte) revenue would be higher for $\tau = 1$, namely $\frac{f_1}{s_1} = \frac{9}{2S} = \frac{4,5}{S}$ which is larger than for the included transactions $\frac{3,5}{S}$, although the average revenue *over the entire* S , that is $\frac{3}{S}$, would be lower with $\tau = 1$.

However if now ($f_1 = 4$; $s_1 = \frac{2S}{3}$), with everything else being the same, then it would be optimal for the miner to confirm only $\tau = 1$, leaving some of the bytes free, as they could accommodate no further transaction. In this case, part of the block capacity will remain empty, *as if* the optimal BSL was smaller than S in this case.

3. Nash Equilibrium of the Transaction Fees Game (TFG)

In the previous section we presented a general approach to discuss the miner's optimal BSL. The approach is based on the assumption that $r(S)$, the expected revenue per confirmed byte, is known by the miner. But how is $r(S)$ determined? The revenue is clearly based on the proposed fees, whose level however emerges as the outcome of the strategic interplay among users wishing to confirm their transactions.

This is why in this section we characterize $R(S)$ as the Nash Equilibrium of what we call the *Transaction Fees Game* (TFG). By this we mean the game played by users, when deciding how much to offer for their transactions to be included in the next block. To keep the analysis simple, though we believe interesting, we shall consider a static version of the game, where fees are proposed independently of each other. Since transactions are executed sequentially, and are publicly observable, a more realistic version of the model should be dynamic, with non-independent choices. Yet, the gain in simplicity of the static game will not compromise the possibility to obtain insightful suggestions. In the next paragraph, we introduce the game notation and fundamentals, reinterpreting the goal of the miner in terms of the proposed fees.

3.1 Revenue of the Successful Miner—Suppose $\tau = 1, \dots, n$ are the outstanding transactions in the *mempool* at time t with s_1, \dots, s_n being their size and f_1, \dots, f_n the associated, proposed, fees. We assume the fee is paid by the payer in the transaction, and in what follows we'll refer to it as his payment proposal for the transaction confirmation. Moreover, let G be the set of all possible non empty $2^n - 1$ groups of pending transactions and $g \in G$ a generic such group. Then, for a given block size limit S , the miner's revenue can be defined as follows:

$$R(S) = Sr(S) = \max_{g \in G} (f(g) = \sum_{i \in g} f_i) \text{ such that } \sum_{i \in g} s_i \leq S \quad (7)$$

That is, given the set of outstanding transactions the successful miner will confirm in the block the subgroup with maximum total fees. Since S is given, an equivalent way to look at (7) is in terms of the average fee per byte, that is

$$r(S) = \max_{g \in G} \frac{f(g)}{S} \text{ such that } \sum_{i \in g} s_i \leq S \quad (8)$$

As discussed in the previous example, the solution to (7)-(8) may satisfy the capacity constraint with strict inequality rather than with equality. That is, for a given BSL, it may be optimal for the miner to leave some bytes empty, as their occupation may imply a decrease in the total fees obtained.

In what follows, for convenience, by $G(S)$ we shall indicate those subgroups g of transactions satisfying the constraint $\sum_{i \in g} s_i \leq S$.

3.2 Nash Equilibrium of the TFG with Complete Information and Uniform Transaction Size—In this section we consider a simple, static, version of TFG, where users submit their fees proposals to confirm their transaction in the next block, that is to be part of the group of transactions solving (7). We assume that offers are submitted independently of each other. That is, the amount chosen as a fee by one user is not known, or if known disregarded, by the other individuals.

As before, assume there are n outstanding transactions and, with no major loss of generality, suppose that $v_1 > v_2 > \dots > v_n$ is the *maximum willingness to pay* (value) for each transaction, ranked from the largest to the smallest. The value of a transaction is defined as the amount of currency which would make the payer indifferent between having it confirmed, paying that fee, or keeping the sum without the transaction being confirmed. We suppose that values are independent of the block size limit S .

Assuming $f_\tau \leq v_\tau \leq b_\tau$, for all $\tau = 1, 2, \dots, n$, then for the payer the payoff obtained from confirmation is given by

$$\Pi_\tau(f_\tau) = \begin{cases} v_\tau - f_\tau & \text{if } \tau \text{ is confirmed in the next block} \\ 0 & \text{if } \tau \text{ is not confirmed in the next block} \end{cases}$$

That is, if confirmation takes place the individual obtains as payoff his willingness to pay for the confirmation v_τ minus the fee, while if confirmation does not take place he would have no

losses due to postponement. Absence of losses may be a somewhat unrealistic assumption, since delayed confirmation is typically undesirable for users. Yet, for the purpose of the paper we find it acceptable.

Before presenting the general results, the following example will help gaining some main insights on how proposed fees could be characterized as a Nash Equilibrium of the game.

Example 2—Consider the simple case in which $s_i = s < S$, that is where the transactions size is the same. Then the number of exchanges that could be inserted in the block is the integer part of $k = \frac{S}{s}$. Suppose BSL is $S = 2$ and $s = 1$. Then the block could include at most two transactions. Further suppose there are four outstanding transactions with the following, uniform, values distribution $v_1 = 4, v_2 = 2,9, v_3 = 1,9, v_4 = 0,9$, and that fees can be expressed with only a single decimal. Then, as we shall argue later, at a *plausible* Nash Equilibrium of the game with complete information the fees profile is

$$f^* = (f_1^* = 2; f_2^* = 2; f_3^* = 1,9; f_4^* \leq 1,9)$$

and the revenue for the successful miner is 4. That is, in equilibrium no confirmed transaction would propose a fee lower than the highest value of the excluded transactions. Moreover, it is optimal for the confirmed transaction fees to be just above it. Therefore, the miner's revenue is positively related to the highest value of the excluded transactions.

It is important to anticipate, and point out, however that the above type of equilibrium is not unique, yet we believe it would be the more plausible. Indeed, the following fees profiles

$$f^{**} = (2 < f_1^{**} \leq 2,9; f_2^{**} = f_1^{**}; f_3^{**} = f_2^{**} - 0,1; f_4^{**} \leq f_2^{**} - 0,1)$$

would also be Nash Equilibria. However at such equilibria, where the miner earns more than with f^* , excluded transactions by players 3 and 4 propose fees which are larger than their values. This is sub-optimal except for when players *are sure* to be excluded from the block, that is they *are sure* not to pay the fees. Indeed, in this case they would make negative profits.

For this reason, we believe such equilibrium fees profiles to be less plausible than f^* , and in what follows we shall focus on the f^* type of equilibrium profile, keeping in mind its non-uniqueness.

Suppose now BSL increases to $S = 3$; then the Nash Equilibrium of the game would be

$$f^* = (f_1^* = 1; f_2^* = 1; f_3^* = 1; f_4^* = 0,9).$$

Therefore, although there is now an additional confirmed transaction the total revenue for the miner reduces to 3. This is lower than with two transactions because the increase in the number of confirmed transactions would not compensate for the reduction in the unit fee.

Likewise if $S = 1$ then the Nash Equilibrium becomes

$$f^* = (f_1^* = 3; f_2^* = 2,9; f_3^* \leq 2,9; f_4^* \leq 2,9)$$

and the revenue, also in this case, lower than with two transactions. So for the miner the optimal BSL, with this distribution of values, is $S = 2$.

Table 1 below summarizes the above examples. Bold figures indicate, respectively, the optimal BSL and the associated miner’s revenue.

Table 1. Block size limit and the miner’s revenue with values $v_1 = 4, v_2 = 2,9, v_3 = 1,9, v_4 = 0,9$.

S	1	2	3	4
$R(S)$	3	4	3	0

However, if the value distribution is highly polarized $v_1 = 10, v_2 = 8,9, v_3 = 1,9, v_4 = 0,9$ then with $S = 1$ the Nash Equilibrium

$$f^* = (f_1^* = 9; f_2^* = 8,9; f_3^* \leq 8,9; f_4^* \leq 8,9)$$

would maximize the miner’s revenue. The point is summarized in Table 2.

Table 2. Block size limit and the miner’s revenue with values $v_1 = 10, v_2 = 8,9, v_3 = 1,9, v_4 = 0,9$.

S	1	2	3	4
$R(S)$	9	4	3	0

Finally, if the value distribution is clustered and has a limited range of variation, such as $v_1 = 1,5, v_2 = 1,3, v_3 = 1,1, v_4 = 0,9$ then with $S = 3$ the Nash Equilibrium

$$f^* = (f_1^* = 1; f_2^* = 1; f_3^* = 1; f_4^* = 0,9)$$

would be the most rewarding for the miner. This point is contained in the following Table 3.

Table 3. Block size limit and the miner’s revenue with values $v_1 = 1,5, v_2 = 1,3, v_3 = 1,1, v_4 = 0,9$.

S	1	2	3	4
$R(S)$	1.4	2.4	3	0

The above example suggests some interesting early insights. The optimal BSL for a miner depends on the distribution of values. If the distribution is polarized with some high values, and the others relatively low, then keeping the BSL small may be optimal. This is because the miner is better off extracting the high willingness to pay of some of the players, rather than trying to include a large number of transactions. However, when the values are relatively close to each other revenue is maximized by including as many transactions as possible in the block, and in so doing increase its size limit. Finally, if the values distribution is relatively uniform, then an *intermediate* BSL is optimal.

Define as $v(g) = \sum_{i \in g} v_i$; then, based on the above assumptions, we can formulate the first result, again referring only to the *plausible* Nash Equilibrium.

Proposition 1 *Suppose the transactions size is symmetric, $s_i = s$, that $f^* = (f_1^*, \dots, f_n^*)$ is a Nash Equilibrium of the game and that $g^* \in G(S)$ is the confirmed group of transactions at such equilibrium. Then $g^* = \{1, 2, \dots, k \leq n\}$ and $v(g^*) \geq f(g^*) = k(v_{k+1} + 0, 1)$. Moreover, all bytes of the block capacity will be occupied if $k = \frac{S}{s}$, while some of the bytes would not be occupied if $k < \frac{S}{s} < k + 1$.*

Proof. Immediate. Since the confirmed transactions will be those with the k highest values, each such transaction will find it optimal to propose the minimum fee needed for confirmation, that is $v_{k+1} + 0, 1$ to avoid ties. Q.E.D.

Consider, for simplicity, the miner’s revenue as kv_{k+1} . The expression clearly shows the trade-off for the miner between including an additional transaction in the block and lowering the single fee. For example, increasing BSL to allow for the confirmation of an additional transaction would be desirable if $kv_{k+1} < (k+1)v_{k+2}$ that is if $k < \frac{v_{k+2}}{v_{k+1} - v_{k+2}}$, namely if k is *small enough*, with the upper limit being defined as a function of the relevant values.

3.3. *The Nash Equilibrium Miner’s Revenue, with Continuous Values and Infinitesimal Transaction Size*—The previous paragraph suggested how the miner’s revenue could be estimated at a Nash Equilibrium of the TFG, where the fee offered by each confirmed transaction is determined by the value of the first excluded transaction. In this section we briefly extend the analysis considering values as continuous random variables, with probability distribution function $P(v)$ and density $P'(v) = p(v)$. Moreover, suppose each transaction occupies an infinitesimal size, and that v stands for the fee proposed by each confirmed exchange, as in Proposition 1. Then, assuming that miners know the values distribution, in equilibrium their revenue $R(v)$ is given by

$$R(v) = v \int_v^\infty p(x)dx = v(1 - P(v)) \quad (9).$$

Hence, the optimal BSL can be defined in terms of the value $v = v^*$ maximizing (9), with the first order condition given by

$$\frac{dR(v)}{dv} = 1 - P(v) - vp(v) = 0 \quad (10).$$

If second order conditions are satisfied then the solution to (10) is given by

$$v^* = \frac{1 - P(v^*)}{p(v^*)} = \frac{1}{h(v^*)} \quad (11)$$

where $h(v) = \frac{p(v)}{1 - P(v)}$ is the hazard rate of the random variable v , namely the probability that the value is in the infinitesimal interval $[v, v + dv]$, given that is no lower than v

Example 3—Suppose v is exponentially distributed with parameter $\lambda > 0$, that is $p(v) = \lambda e^{-\lambda v}$ with $v \geq 0$. Then $h(v) = \lambda$ and $v^* = \frac{1}{\lambda} = Ev$. Therefore, in this case miners should set BSL to include all transactions above the expected value of v , confirming a share of the pending transactions equal to about $\int_{\frac{1}{\lambda}}^{\infty} \lambda e^{-\lambda v} dv = 1 - \frac{1}{e}$.

If instead v is uniformly distributed over the closed interval $[a, b]$ with $0 < a < b$, it is $p(v) = \frac{1}{(b-a)}$ and $P(v) = \frac{v-a}{b-a}$. Then $v^* = b - v^*$ and so $v^* = \text{Max}(a, \frac{b}{2})$.

3.4. The Nash Equilibrium of the TFG Game with Complete Information and General Transaction Size—In this paragraph we discuss the TFG in the general case, where transactions can have different sizes. Based on the above assumptions, we can formulate the following result, again referring only to the *plausible* Nash Equilibrium.

Proposition 2 *Suppose $f^* = (f_1^*, \dots, f_n^*)$ is a Nash Equilibrium of the TFG game and that $g^* \in G(S)$ is the confirmed group of transactions at such equilibrium. Then (i) $v(g^*) \geq f(g^*) \geq \max_{\{g \in G | g \cap g^* = \emptyset\}} v(g)$ and (ii) for all transactions pairs i and j of the same size, with $i \in g^*$ and $j \notin g^*$, it is $f_i^* \geq v_j^*$.*

Proof. (i) Suppose g^* is confirmed at a Nash Equilibrium and that for some $g \cap g^* = \emptyset$ it is $v(g^*) < v(g)$. Then, there must be a subset of transactions $g' \subseteq g$ for which it is possible to optimally increase the proposed fee in such a way that $f(g^*) < f(g) < v(g)$ so that g would be confirmed, contradicting the initial assumption. (ii) An argument similar to (i) applies. Q.E.D.

The above proposition shows that with general transactions size, at a Nash Equilibrium the confirmed set of transactions fees must satisfy two types of constraints. The first is an *aggregate* type of constraint $f(g^*) \geq \max_{\{g \in G | g \cap g^* = \emptyset\}} v(g)$ and the second is an *individual transaction* type of constraint. Their meaning, and why both such constraints are needed, is illustrated by the following example

Example 4—Start considering $n = 4$, $S = 4$ and $v_1 = 4, v_2 = 0,9, v_3 = 1,9, v_4 = 1,9$. Moreover $s_1 = 3, s_2 = 1, s_3 = 2, s_4 = 2$. Then, as we shall argue later, at all *plausible* Nash Equilibria of the game the confirmed transactions are $g^* = \{1,2\}$ and $f(g^*) = 3.9 \geq 3,8 = \max_{\{g \in G | g \cap g^* = \emptyset\}} v(g)$. The following two points are worth noticing:

i) Given BSL, the only groups of transactions that could be confirmed are $\{1,2\}, \{3,4\}, \{2,3\}, \{2,4\}, \{1\}, \{2\}, \{3\}, \{4\}$ where, clearly, in equilibrium $\{1,2\}$ is the only one inducing the highest revenue for the miner.

ii) There is a multiplicity of *plausible* Nash Equilibria, all being characterized by $f(g^*) = 3.9$. For example,

$$f^* = (f_1^* = 3.9; f_2^* = 0; f_3^* = 1,9; f_4^* = 1,9)$$

and

$$f^{**} = (f_1^{**} = 3; f_2^* = 0.9; f_2^* = 1,9; f_4^* = 1,9)$$

are two of them, where the second transaction either pays no fees or its maximum willingness to pay. Moreover, there could also be equilibria with *intermediate payments* such as

$$f^{***} = (f_1^{***} = 3.5; f_2^* = 0.4; f_2^* = 1,9; f_4^* = 1,9).$$

Which of these equilibria would prevail is hard to say *a priori*, and would very much depend on if and how individuals could coordinate their fees proposals. In any case, what really counts for the miner is that $f(g^*) = 3.9$, while for the users paying to confirm their transactions, spending less is clearly preferable than more. Hence, the first user would certainly prefer f^* to f^{**} while the second user would have opposite preferences.

In the above example, the only constraint satisfied at the equilibrium is the *aggregate* constraint, that is $f(g^*) \geq 3.9$, but confirmed transactions $\{1,2\}$ do not have to satisfy *individual* constraints since the excluded exchanges $\{3,4\}$ each have different sizes from the included ones.

However, suppose now everything is the same except for an additional fifth transaction with value $v_5 = 0,4$ and size $s_5 = 1$. It is easy to see that also in this case $g^* = \{1,2\}$ and $f(g^*) = 3.9$, but now it must be $f_2^* \geq 0.5 > 0.4 = v_5$ because otherwise the second transaction could be excluded by the fifth, which is not optimal for the former. Hence in this case both the aggregate constraint and this last individual constraint are at work in establishing Nash Equilibria.

The above result generalizes Proposition 1 in that the second highest group value $\max_{\{g \in G | g \cap g^* = \emptyset\}} v(g)$ would serve as reference for the miner's revenue. Therefore, also in a more general scenario, the miner's revenue is determined by the distribution of values and transactions size. Again, broadly speaking, if the distribution is *skewed* towards high values then it would be more profitable for the miner to reduce BSL, while if more uniformly distributed then it would be preferable for the miner to enlarge BSL.

4. Discussion and Conclusions

In the paper we discussed BSL and transaction fees in Bitcoin from multiple perspectives. As for blocks of confirmed transactions we investigated the conditions for their size limit to maximize the miner's revenue. The analysis suggests that for the successful miner, the optimal BSL is determined by how much users are willing to pay to confirm their exchanges, which we defined as the *value* of a transaction. Our findings indicate that, whenever possible, at a Nash Equilibrium of the game, fees of confirmed transactions tend to be just above the highest value of the excluded transactions. The game is modelled as an auction, where the successful miner is selling shares of the block space for transactions confirmation, and users bid their fees to occupy such space.

However, the model and related insights, raise few operational and methodological questions.

- (i) The first is how can individuals estimate transaction values? If it is reasonable to assume that users know their own value, it is much less realistic to suppose that they know each others' values and, moreover, that also miners know them. However, the time series of f_τ and/or b_τ could provide interesting information on the distribution of values. Indeed, while we suppose that values do not depend on the BSL, it is reasonable to think that they are positively related to both the proposed fees and the sum exchanged. The observed distributions of fees and/or exchanged bitcoins, individually or in combination, could be used to estimate the value distribution. Below we sketch how these considerations would also be helpful to make the methodology, to determine the optimal BSL, operational directly in terms of the fee distribution, rather than of the values distribution.
- (ii) In Bitcoin, BSL is known in advance to users, and based on such knowledge they propose their fees. This is *as if* miners and users would play a dynamic game, with the former agreeing on BSL and committing to it, and the latter reacting to BSL by proposing their fees. With such an arrangement, agreeing on a limitless block size, so that all outstanding transactions would be confirmed, would lead to zero fees submitted. An alternative could be to agree that all positive fees, above a specified threshold, would be included. In this case, choice of the minimum fee to propose can follow the same methodology as in Section 3.3. In particular, suppose $\Pi(f)$ is the transaction fees distribution function, estimated from the observed time series of fees, and $\pi(f) = \Pi'(f)$ the estimated density. Then the revenue maximising minimum fee f_m would be found by solving

$$\max_{f_m} f_m \int_{f_m}^{\infty} \pi(x) dx = f_m (1 - \Pi(f_m)) \quad (16).$$

Assuming that first order conditions identify the solution to (16), it follows that

$$f_m = \frac{1 - \Pi(f_m)}{\pi(f_m)} = \frac{1}{h(f_m)}$$

where, here too, $h(f_m)$ is the hazard rate of the fee distribution. That is, observed transaction fees would be used as a proxy for values. For instance, if in the past miners observed that fees are uniformly distributed between 0 and 10 bitcoins, then a way to proceed for them would be to agree on confirming all pending transactions proposing at least 5 bitcoins as fees.

Clearly, so far available data on transaction fees have been conditional on the BSL decided by the community, and kept as fixed until now. This could bias the fee distribution estimate for a BSL that may change in the future. However, use of available data may be a first step towards the computation of an optimal BSL for miners.

We conclude observing that an optimal BSL for the users is likely to be different from that for the miners, a point which may suggest directions for further research.

Acknowledgement

The author thanks Bitcoin Unlimited for financial support and Peter Rizun for constructive comments on a previous version of the paper. I would also like to thank the Editor and three anonymous referees whose comments much improved the paper. The paper reflects the view of the author only.

Conflict of Interest

Though the project has been funded by Bitcoin Unlimited, all views expressed in the paper are the author's.

Notes and References

¹ Houy, N. “The Bitcoin Mining Game.” *Ledger* 1 53-68 (2016)

<https://doi.org/10.5195/ledger.2016.13>.

² Dimitri, N. “Bitcoin Mining as a Contest.” *Ledger* 2 31-37 (2017)

<https://doi.org/10.5195/ledger.2017.96>.

³ Huberman, G., Leshno, J., Moallemi, C. “An Economic Analysis of the Bitcoin Payment System.” Columbia Business School Research Paper 17-92 (2019)

<https://dx.doi.org/10.2139/ssrn.3025604>.

⁴ Houy, N. “The Economics of Bitcoin Transaction Fees.” GATE Working Paper 1407, 2014 (2014) <https://dx.doi.org/10.2139/ssrn.2400519>.

⁵ Kaşkaloğlu, K. “Near Zero Bitcoin Transaction Fees Cannot Last Forever,” in SDIWC, *The International Conference on Digital Security and Forensics (DigitalSec2014)* 91-99 (2014)

<http://sdiwc.net/digital-library/near-zero-bitcoin-transaction-fees-cannot-last-forever.html>.

⁶ Moser, M., Bohme, R. “Trends, Tips, Tolls: A Longitudinal Study of Bitcoin Transaction Fees,” in M. Brenner, N. Christin, B. Johnson, K. Rohloff (eds) *Financial Cryptography and Data Security FC 2015 International Workshops, BITCOIN, WAHC, and Wearable, San Juan, Puerto Rico, January 30, 2015, Revised Selected Papers* Berlin: Springer 19-33 (2015)

https://doi.org/10.1007/978-3-662-48051-9_2.

⁷ Rizun, P. “A Transaction Fee Market Exists Without a Block Size Limit.” Bitcoin Unlimited Block Size Debate Working Paper (2015)

<https://www.bitcoinunlimited.info/resources/feemarket.pdf>.

⁸ Easley, D., O’Hara, M., Basu, S. “From Mining to Markets: The Evolution of Bitcoin Transaction Fees.” *Journal of Financial Economics* (available online) (2019)

<https://doi.org/10.1016/j.jfineco.2019.03.004>.

⁹ Lavi, R., Sattath, O., Zohar, A. “Redesigning Bitcoin’s Fee Market.” *arXiv* (2017) (accessed 27 May 2019) <https://arxiv.org/abs/1709.08881>.

¹⁰ Chepurnoy, A., Kharin, V., Meshkov, D. “A Systematic Approach To Cryptocurrency Fees,” in A. Zohar, I. Eyal, V. Teague, J. Clark, A. Bracciali, F. Pintore, M. Sala (Eds.) *Financial Cryptography and Data Security FC 2018 International Workshops, BITCOIN, VOTING, and WTSC, Nieuwpoort, Curaçao, March 2, 2018, Revised Selected Papers* Berlin: Springer 19-30 (2019) https://doi.org/10.1007/978-3-662-58820-8_2.

¹¹ Wilson, R. “Auctions of Shares.” *Quarterly Journal of Economics* **93.4** 675-689 (1979) <https://doi.org/10.2307/1884475>.

¹² Bonneau, J. “Bitcoin Mining Is NP-Hard.” *Freedom to Tinker* (27 October 2014) <https://freedom-to-tinker.com/2014/10/27/bitcoin-mining-is-np-hard/>.



Articles in this journal are licensed under a Creative Commons Attribution 4.0 License.



Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.