



Nuffield  
College  
UNIVERSITY OF OXFORD

# ECONOMICS DISCUSSION PAPERS


---

Paper No: 2023-W04

An evaluation of the Bank of England's ILTR operations:  
Comparing the product-mix auction to alternatives

By Julia Giese and Charlotte Grace

December 2023



# An evaluation of the Bank of England’s ILTR operations: Comparing the product-mix auction to alternatives

Julia Giese\* and Charlotte Grace†

October 2023

## Abstract

We compare the product-mix auction (PMA)—the mechanism used by the Bank of England (BoE) for its Indexed Long-Term Repo (ILTR) operations—to simpler alternative auction designs, namely a pair of separate simultaneous auctions, and a “reference price auction”. Using data from the auctions held in June 2010 to January 2014, we find that the PMA increased welfare (defined by the difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept) by approximately 50%, or 2 basis points per loan, relative to these alternatives. We would expect larger welfare gains in a less stable period than the period studied, and simulations confirm this. Broader benefits of the auctions of reducing systemic risk, while mitigating moral hazard, informing the BoE about stress in the market, and communicating the “correct” price to the market, are taken into account in our approach, to the extent that the BoE’s supply curve internalises some of these externalities. We also find that the PMA always gave the BoE more (or occasionally the same) surplus and revenue relative to if one of the alternative designs had been used. However, the effect of the PMA on aggregate bidder surplus was ambiguous. The latter result may be a property of the period studied, and of the fact that there were only two sets of eligible collateral in this period.

Key words: Product mix auction, auction design, central bank liquidity provision.

JEL classification: D44 (auctions), E58 (central banks and their policies).

---

\*Bank of England. Email: [julia.giese@bankofengland.co.uk](mailto:julia.giese@bankofengland.co.uk)

†Nuffield College, University of Oxford. Email: [charlotte.grace@nuffield.ox.ac.uk](mailto:charlotte.grace@nuffield.ox.ac.uk)

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful to Charlotte Barton, Rand Fakhoury, Andrew Hauser, Tom Horn, Dennis Jeevarajasingham, Paul Klemperer, Rhys Phillips, and an anonymous referee.

## 1 Introduction

The Bank of England (BoE) introduced the Indexed Long-Term Repo (ILTR) auctions in June 2010 to replace the long-term reverse repo operations with a mechanism that would efficiently respond to “the new demands for liquidity insurance that [the financial crisis] engendered” (Fisher, 2011, p.15). The BoE used a novel design in which the quantity of funds lent against different types of collateral responded flexibly to the economic environment. The main objective of the design was to reduce systemic risk while avoiding moral hazard, and it also improved efficiency directly.

The ILTR auctions were designed to be conducted as product-mix auctions (PMA), which are single-round auctions designed to address the specific challenges of lending funds (i.e. central bank reserves) to financial institutions against multiple types of collateral. The maximum supply of funds is fixed, but the quantity of funds lent against collateral within a particular set varies depending on the bids received and the BoE’s own supply preferences.<sup>1,2</sup>

This paper assesses both the benefits of the design of the ILTR in terms of welfare gains, and the distribution of benefits between the BoE and bidders.

On the demand side, we assume that participants’ bids express the prices they are willing to pay.<sup>3</sup> We measure the surplus of each participant by the difference between their bid and the price they pay (and sum this across bidders to get total bidder surplus). This does not take account of any externalities that result from meeting a participant’s liquidity need and from charging it the correct price for doing so, but it equals their individual surplus.

On the supply side, our analysis relies on the mechanism being designed to maximise the overall benefits to society, that is, the BoE’s own supply preferences express the “correct” combinations of prices and quantities. Any broader benefits (see below) of allocating funds to bidders who are willing to pay a price above the BoE’s supply curves are therefore taken

---

<sup>1</sup>Since 2014, the BoE does not commit to a maximum supply, and the total quantity of funds in the ILTR operations also responds to the bids and BoE’s preferences. We therefore focus on the period June 2010 – January 2014 in which the same design was used for all auctions (except for some minor pricing changes, detailed in Section 2).

<sup>2</sup>Until October 2013, Level A and Level B collaterals were referred to as Narrow and Wide collaterals, respectively. They have also been respectively referred to as Strong and Weak collaterals. The Level A collateral set includes gilts, sterling Treasury bills and certain sovereign and central bank debt, whereas the Level B set includes less liquid sovereign debt and certain asset-backed securities. The BoE’s current auctions also allow bids on (and generates a separate price for) loans of funds secured by collateral in a Level C set, which currently includes less liquid securities and portfolios of loans. For a full list of eligible collateral, see <https://www.bankofengland.co.uk/markets/eligible-collateral>.

<sup>3</sup>Grace (in prep., a) finds that this appears to be a reasonable approximation in the ILTR auctions.

into account to the extent that the supply curves internalise some of these externalities.

Importantly, we find that the welfare gain increases significantly when simulating a stress event, i.e. the benefits are more readily observed in times of higher liquidity demand.

To evaluate the PMA design, we compare alternative auction rules to the actual ILTR auctions that were held monthly in June 2010 – January 2014. The two main comparators are (i) separate simultaneous auctions (SSA) for loans secured by the different sets of collateral, each with a fixed supply of funds and (ii) a single-round “reference price auction” (RPA), in which the auctioneer fixes a pair of notional prices for the two collateral sets, i.e. “reference prices”, and the highest bids relative to their reference prices are accepted. Setting reference prices in the RPA is equivalent to fixing the difference between the market clearing interest rates for the collateral sets in advance. We choose these two comparators as simpler reasonable alternatives, used either historically or by other central banks. For example, the BoE’s gilt purchase operations, introduced in 2009, are run as RPAs, and the BoE’s Corporate Bond Purchase Scheme, introduced in 2016, is a variant of an RPA. In our setting, we compare net welfare, bidder surplus, BoE revenue, and BoE surplus across auctions, for various fixed quantities (in the SSA), and fixed reference prices (hence, fixed price differences) (in the RPA). We also compare the outcomes of the PMA to separate sequential auctions and to posting fixed spreads over Bank Rate (with rationing if necessary) to borrow against the two collateral types.

The PMA jointly determines the quantities of funds lent against Level A and Level B collateral. It finds the competitive equilibrium, assuming that bids express bidders’ true valuations for loans, i.e. bids are “truthful”, and assuming that the BoE’s supply curves express the minimum prices that it is willing to accept to supply loans secured by the collateral sets;<sup>4</sup> the difference between the market clearing interest rates (“prices”) for funds lent against Level A and Level B collateral varies across auctions to implement these allocations. Under these assumptions, the PMA therefore maximises welfare (as conventionally measured as the sum of the BoE’s and bidders’ surpluses).<sup>5</sup> The alternative auction designs may fail to maximise welfare because of the constraints they introduce: the SSA and RPA do not permit quantities and the price difference, respectively, to adjust to the participants’ preferences. Unlike the PMA, neither of these natural alternatives therefore allow a substantial increase

---

<sup>4</sup>Full details in Section 2.

<sup>5</sup>This is the case providing the clearing prices for the two goods are chosen to be the competitive equilibrium prices. The ILTR rules in the earlier part of the sample period permitted the clearing prices to occasionally differ from the competitive equilibrium, but this made little difference to the auction outcomes, and we describe the PMA as “welfare-maximising” throughout. See Appendix A for details.

in the allocation of funds against the less liquid Level B collateral in the event of stressed conditions, at prices the BoE is willing to accept. In Section 7, we consider alternative objectives to welfare, including maximising BoE revenue or surplus.

Our main finding is that the flexibility in allocation of the PMA leads to welfare gains of approximately 50%, or 2 basis points per loan, as measured by the difference between the spreads that participants were willing to pay and the spreads that the BoE was willing to accept, relative to these two alternative designs. This is the case even though our period of analysis, June 2010 – January 2014, was a period without significant stress. Indeed, even with a hypothetical alternative objective of maximising BoE revenue or BoE surplus, rather than welfare, the PMA outperforms all possible RPAs and almost all possible SSAs. The benefits result from the automatic adjustment of the PMA allocation to the bids received and to the BoE’s preferences between the two collateral sets. While the bidders were permitted to flexibly represent their preferences through “paired bids” for funds lent against Level A or Level B collateral over this period, this feature was rarely used.<sup>6</sup>

A second finding is that, in the period studied, the PMA always gave the BoE more (or occasionally the same) surplus and revenue relative to if the BoE had run any possible SSA or RPA, but the effect on the bidders, in aggregate, was ambiguous. It is trivial that the PMA must (weakly) increase aggregate surplus. However, it was slightly surprising that bidders, on average, shared in the gains only when the fixed share to Level B used in the SSA was low, or when the difference in fixed reference prices between A and B used in the RPA was high, so that the PMA typically lent more against Level B relative to the alternative. These results may be a property of the period studied, and of the fact that there were only two sets of eligible collateral, as we discuss below. We expect, and confirm in simulations, that the welfare gain and benefits to the BoE of the PMA would be qualitatively similar, but quantitatively larger in absolute terms, in a less stable period than the period studied. This reflects the scaling up of the benefits of the PMA’s flexibility in the allocation and sensitivity to the BoE’s preferences in contrast to the comparators. We also expect, as we explain and confirm in simulations below, that the specific features of the sets of eligible collateral mean that the PMA would benefit bidders over a larger range of SSAs, but a smaller range of RPAs, in a less stable period.

Importantly, none of our calculations take any explicit account of the broader benefits of

---

<sup>6</sup>Each bid of this kind specifies the quantity of loans demanded by the bidder and a bid price for each of the two collaterals. The bidder wins only if they would obtain positive surplus, and wins at most the quantity demanded of funds, borrowed against the collateral which gives them the highest surplus given their bids.

the PMA of reducing systemic risk while avoiding moral hazard (as well as informing the BoE about stress in the market and communicating the “correct” price to the market). The BoE wanted to permit banks to repo using collateral that was lower quality than the Level A collateral it had traditionally accepted, but only at prices which reflected the relative liquidity of the collateral assets. Moreover, the quantity of less liquid (Level B) collateral it was willing to accept was an increasing function of market stress. Clearly neither the SSA nor the RPA could achieve as fine-tuned a response as the PMA to market conditions. And so, because it does not capture all the broader benefits, our measure can be interpreted as a conservative estimate of the overall welfare gains of the PMA.

Section 2 provides background for the ILTR auctions and Section 3 describes our method. We compare the PMA to the SSA in Section 4 and discuss the impact of including reserve prices in the SSA, and of running the separate auctions sequentially. Section 5 compares the PMA to the RPA and to posting fixed prices. We expect the surplus comparisons would be more extreme in periods of stress, and Section 6 shows simulations that confirm this. Section 7 shows that alternative objectives to welfare-maximisation have only limited qualitative impact on the results. Section 8 describes the broader benefits of the PMA. Section 9 concludes.

## 1.1 Welfare results

**SSA** Had the BoE been able to forecast the optimal fixed quantities of funds to lend against Level A and Level B collateral separately, the SSA would have been almost as efficient as the PMA over the period studied. However, it seems unrealistic that the BoE could have identified this specific, best-performing SSA prior to the auctions’ implementation. It introduced the PMA to overcome many operational policy questions relating to the existing LTR operations. In particular, the size of the operations and the appropriate spread between lending against Level A and Level B collateral were very unclear. The large increase in demand for funds in 2008 – 2009 added to this uncertainty. If the BoE had over- or underestimated the optimal split of its total allocation of funds between the two separate auctions, the welfare loss incurred could have been substantial.

Importantly, the PMA allows the quantity allocated against Level B collateral to respond to stress. The composition of demand across Level A and Level B collateral in June 2010 – January 2014 turned out to be relatively stable. Had it not been, then any fixed split between the two collaterals in an SSA would have resulted in a significant welfare loss. For example, an unexpected increase in stress would have caused a potentially dramatic shift in

demand towards lending against Level B collateral.<sup>7</sup> Focusing on the period of June 2010 – January 2014 risks underestimating the advantages of the PMA relative to the SSA. When designing the auction, the BoE was seeking a mechanism that accommodated demand in periods of acute stress, rather than the relatively stable demand that is observed.

In particular, given that the objective of the ILTR was to lend funds against Level B collateral in stressed circumstances, any fixed split between Level A and Level B collaterals, chosen by the BoE, would presumably have allocated a significant proportion against Level B collateral. Our results show that any fixed allocation against Level B exceeding 50% would have generated a welfare loss exceeding 2 basis points (bps) per loan of the quantity of funds in the ILTR operations and 35% of the welfare generated by the PMA.

**RPA** The difference in clearing prices between funds lent against Level A and Level B collaterals in the RPA equals the difference in their “reference prices”, which are fixed in advance. There was enough variation in the optimal price difference between A and B across auctions that no single pair of reference prices came very close to maximising welfare in the period studied.

Across this period, the best-performing RPA would have generated an average welfare loss of 0.6bps and 0.8bps respectively per loan over a 3- and 6-month term, equivalent to 10% and 14% respectively of the average welfare generated by the PMA. Of course, this presumes the BoE would be able to choose the optimal price difference between collateral sets in advance. A reference price of B of 50bps and reference price of A normalised to zero (corresponding to the reserve prices set for Level B and Level A collateral in the auctions from October 2008 up until the June 2010 start of the PMA) would have almost maximised welfare for loans over a 6-month term, yet the average loss would have been around 3bps for loans over a 3-month term (equal to 53% of PMA welfare). In contrast, a reference price of B of 0bps (as in auctions prior to October 2008) would have generated an average loss for both the 3- and 6-month terms, of 2 and 5.3bps (32% and 71% of PMA welfare) respectively.

## 1.2 Benefits to bidders versus benefits to the BoE

The BoE always, and often significantly, benefited in both revenue and surplus (as measured by the difference between the market clearing prices for the loans and the prices the BoE

---

<sup>7</sup>This was observed in March 2020, in which the financial stress associated with COVID-19 led to a sharp increase in bids for funds lent against Level B collateral and especially against Level C, an additional, less liquid, collateral set introduced to the ILTR in 2014, after the sample period.

was willing to accept) from the PMA auction, relative to any possible SSA or RPA.<sup>8,9</sup>

However, perhaps surprisingly, bidders benefitted from the PMA relative to only some SSAs and RPAs.<sup>10</sup> Several features seem to have contributed to this.

First, the nature of demand for loans secured by Level A collateral is that it is always relatively flat at a price close to zero (and there is a reserve price of zero), so bidders can never make substantial gains or losses on Level A, whatever the auction form.

Second, crucially, in the time period in question, there was never very significant stress. As a result there was neither ever very large demand for loans secured by Level B, nor ever very large net surplus to bidders from borrowing against Level B (as bids were rarely much above the relevant point of the BoE's supply curve for Level B). Bidders would therefore have lost relatively little if the quantity lent against B had been below the welfare-maximising quantity, but would have gained substantially from very cheap borrowing against Level B if the quantity had been excessively large. Almost every SSA or RPA, fixed across the period, sometimes lent too little against B, and at other times too much, relative to the welfare-maximising PMA. The inefficiencies of the alternative designs benefitted the bidders in this respect.

A third feature was that bidders rarely expressed their preferences across the different collaterals, despite being able to do so via paired bidding.

Absent any one of these three features, the bidders might have benefited more from the use of the PMA.

If demand for loans secured by Level B collateral were more variable and sometimes much greater, the welfare benefits of the PMA relative to the SSA would likely have provided benefits to bidders, as well as the BoE. This is because the PMA would increase the alloca-

---

<sup>8</sup>There was one tiny exception over the period studied: running an RPA for 6-month loans with a fixed price difference of 50bps increases average BoE surplus by 0.01bps relative to the actual PMA.

<sup>9</sup>For any one auction, the PMA is, of course, identical to the SSA and RPA that are optimally calibrated for that auction, with the exception that the RPA may ration the bids incorrectly at the marginal prices. If demand is fairly stable, the PMA and an SSA that is optimally calibrated to the whole set of auctions are close to identical, but any given RPA may still lead to welfare losses, as discussed in Section 5.

<sup>10</sup>Grace (in prep., b) shows that the PMA increases both bidder and auctioneer surplus relative to the best-performing SSA under reasonable conditions in a theoretical model in which bidders face correlated shocks to their values. She shows that these predictions are consistent with the estimated effect of using a PMA in Mexican Treasury auctions. These results suggest that both the bidders and the BoE would gain from the PMA. The smaller benefits we find for the bidders may be because shocks to bidders' values are partly idiosyncratic. Lings (2013) demonstrates that this reduces the likelihood that bidders gain from the PMA.

tion of the relatively valuable loans against Level B in response to the relatively high bids. However, an RPA with a low price difference between Levels B and A would continue to largely prioritise bids on B, so that the PMA would not necessarily generate greater benefits for bidders. Some experimentation with different possible demands confirmed this claim (see Section 6).

If there was more than one collateral for which the demand curve was not flat, the PMA might provide more benefits to bidders. In particular, in a three-collateral implementation (Levels A, B and C, with C being the least liquid), such as the BoE currently uses, the bidders, as well as the BoE, might gain from the PMA, if the demand curves for two of the collaterals—presumably Levels B and C—are not flat.<sup>11</sup> This could be explored in further work.

And if bidders used the paired bids facility, they could directly gain from the use of the PMA (without loss to the BoE), in cases where they have the ability to substitute collaterals. Perhaps they did not have this ability, but it may be that they would in the future.<sup>12</sup> Alternatively bidders may not have been very sophisticated.<sup>13</sup> Again, their understanding could improve in the future.

## 2 Background

We study the ILTR auctions, conducted monthly by the Bank of England, over the period June 2010 – January 2014. Table 1 presents summary statistics.

The auction is a product-mix auction (PMA), described in Section 2 of Klemperer (2010, 2018).<sup>14</sup> In each auction, loans of central bank reserves for either a 3- or 6-month term are

---

<sup>11</sup>Bidders (as well as the BoE) would likely gain from the PMA relative to either using the SSA with a high share to Level C relative to Level B, or the RPA with a low price difference between Level C and Level B.

<sup>12</sup>As BoE Executive Director, Andrew Hauser, wrote (a little after the BoE decided it would stop using the paired-bids facility from February 2014, although its current programme is designed to accommodate paired bids), “higher collateral prices will give financial firms a strong incentive to invest in technologies allowing them to optimise their use of collateral. Indeed, big investment programmes are already underway in many firms and infrastructure providers, to ensure that firms have real-time information on the collateral they have available globally across all their business lines, that the collateral they deliver is cost effective, and that the cost of delivering (and financing) that collateral is factored into their risk and business decisions. These programmes involve sometimes relatively advanced technology; indeed, as some of our contacts remark, somewhat alarmed, ‘for the first time in living memory, pointy heads are sitting on the repo desk’.” <https://www.bankofengland.co.uk/-/media/boe/files/speech/2013/the-future-of-repo-too-much-or-too-little>

<sup>13</sup>Preliminary results (Grace, in prep., a) suggest that bidders were not too sophisticated in their use of the PMA, more generally — the good news is that this also means that strategic behaviour was unlikely (see Section 3.1).

<sup>14</sup>Klemperer (2010) is a minor revision of Klemperer (2008). Klemperer (2018) adds additional material

Table 1: Summary statistics for ILTR auctions, June 2010 – January 2014

		3-month term		6-month term	
		Level A	Level B	Level A	Level B
Quantity of funds demanded (weighted, %)	Mean	79.69	20.31	66.44	33.56
	Std. Dev.	17.01	17.01	27.15	27.15
Quantity of funds allocated (weighted, %)	Mean	83.28	16.72	67.46	32.54
	Std. Dev.	13.41	13.41	24.56	24.56
Price paid (weighted average, bps)	Mean	1.17	22.12	0.9	38.91
	Std. Dev.	1.74	6.8	0.86	16.7
Bid price (weighted, bps)	Mean	2.5	22.98	3.06	35.87
	Std. Dev.	2.85	10.28	4.31	17.14

auctioned at a spread over Bank Rate.<sup>15</sup>

Within each auction, bidders may borrow the reserves against one of two types of collateral: Level A, including gilts, sterling Treasury bills and certain sovereign and central bank debt, and Level B, including less liquid sovereign debt and certain asset-backed securities.

**Bids** Bidders can submit any number of sealed bids, each of which can be either of two types:

1. Specific bids: Each bid of this kind specifies the type of collateral the bidder will provide, the quantity of funds that they demand, and the spread that they are willing to pay for the funds borrowed against the specified type of collateral. For example, a bid may specify demand of £50 million for a spread of 2bps against Level B.
2. Paired bids: Each bid specifies the quantity of funds the bidder demands, and the spread that they are willing to pay for the funds borrowed against each of the Level A and B collaterals. For each bid, they may be allocated the demanded funds against *either* Level A *or* Level B collateral.<sup>16</sup> For example, a bid may specify demand of £50 million for a spread of 2bps against A or 25bps against B. The bid will be allocated up to £50 million against *either* Level A *or* Level B.

**Paired bids** In practice, bidders rarely submit paired bids so we drop these from the

to Klemperer (2010), but their Section 2s are identical.

<sup>15</sup>Loans in the current ILTR auctions are for 6-month terms. The Short-Term Repo (STR) facility (announced in August 2022) will run alongside the ILTR auctions and offer loans with one-week maturity against Level A collateral at a fixed at Bank Rate (with no maximum supply).

<sup>16</sup>If both paired bid prices are higher than the respective clearing prices, the bid is allocated to the good that generates the highest surplus to the bidder (determined by the difference between the bid and clearing price).

sample. This change is insubstantial.

The minimum bid size is £5 million, with increments of £1 million, and the minimum unit of allocation is £100,000. There is no obligation for registered bidders to participate in an auction, and bidders can bid up to 30% of the maximum supply (£1.5 billion and £750 million in the 3- and 6-month term auctions respectively).

**Supply** Prior to any one auction, the BoE commits to two supply curves, which are illustrated in Figure 1:

1. A maximum supply: The BoE announces a maximum supply made available across both Level A and B collaterals. The maximum supplies were £5 billion in each of the 3-month term auctions and £2.5 billion in each of the 6-month term auctions.<sup>17</sup>
2. A “relative supply” curve: The BoE commits to a privately known “relative supply” curve. This represents the minimum price that the BoE is willing to accept on the margin to supply funds against less liquid Level B collateral relative to Level A collateral.<sup>18</sup> The BoE’s supply preferences are therefore measured in terms of the difference between the spread for Level B collateral and the spread for Level A collateral, as a function of the percentage of the maximum supply allocated against Level B collateral. This influences the welfare-maximising allocation across the two collateral sets.


We refer to funds borrowed against each type of collateral as goods  $A$  and  $B$  respectively. The market clearing spreads over Bank Rate for loans of funds secured by each set of collateral are referred to as the clearing prices for goods  $A$  and  $B$ .

---

<sup>17</sup>Since 2014, the BoE has committed to an upward-sloping “total supply”, which is a function of the total quantity allocated across all collateral sets, rather than a maximum supply. The maximum supply in our context is equivalent to a total supply curve which is flat at zero up to the maximum supply, at which point it becomes vertical.

<sup>18</sup>The difference in the prices that the BoE is willing to accept to lend against Level B and A collaterals reflects differences in liquidity. (The BoE imposes a haircut to each collateral asset to account as far as possible for differences in risk, see de Roure and McLaren (2021).)

Figure 1: Competitive equilibrium prices,  $(p_A, p_B)$ , and allocations,  $(q_A, q_B)$ , in the PMA. The welfare gain from allocating a marginal unit is equal across the two goods, i.e.  $\Delta_A = \Delta_B$



**Market clearing rules** The PMA is a uniform-price auction, so that all winning bids pay the single market clearing price for the good they win. As described in Klemperer (2008), the original proposal was that these two prices (one each for goods  $A$  and  $B$ ) were to be chosen to be competitive equilibrium prices. This would mean that the auction would maximise social welfare (the sum of BoE and bidder surplus), assuming bidders bid their actual valuations, which is a reasonable assumption.<sup>19</sup> To implement this, the PMA finds the quantities such that the difference between the clearing prices for goods  $B$  and  $A$  is equal to the difference in the prices the BoE is willing to accept on the margin between  $B$  to  $A$ . Figure 1 shows the competitive equilibrium prices, and the corresponding quantities allocated, for an illustrative realisation of bids. It is easy to see from the figure that either increasing or reducing  $q_A$  (and changing  $q_B$  correspondingly) would reduce total welfare.<sup>20</sup> See Section 3.1 and Appendix A for discussion.

However, even where the competitive equilibrium prices are unique (they generally are—see Appendix A.1) there is often a range of prices at which demand (although not the BoE’s

<sup>19</sup>At least for bidders who are not bidding for too large a fraction of the amount offered, bidding truthfully is (close to) optimal.

<sup>20</sup>Total welfare is measured as the sum of the area under the aggregate demand curve for good  $A$  up to  $q_A$  and the area between the aggregate demand curve for good  $B$  and the relative supply curve up to  $q_B$  (see Section 3.3).

Table 2: PMA market clearing rules for the ILTR auctions, June 2010 – January 2014

Operation Date	Price Determination		Relative Supply Curve	
	Good (Level) A	Good (Level) B	3-month term	6-month term
Jun 2010 – Apr 2011	Lowest-winner pricing	Lowest-winner pricing	Starts at 5bps	Starts at 5bps
May 2011 – Aug 2011	Lowest-winner pricing	Lowest-winner pricing	Starts at 5bps	Starts at 15bps
Sep 2011 – Apr 2012	Lowest-winner pricing	Highest-loser pricing	Starts at 5bps	Starts at 15bps
May 2012 – Jan 2014	Highest-loser pricing	Highest-loser pricing	Starts at 5bps	Starts at 15bps

offered supply) matches the competitive equilibrium quantities. Appendix A describes the precise price determination rules, and Table 2 summarises them as they vary over the period. On both  $A$  and  $B$ , the rules changed from “lowest-winner pricing” (LWP) to “highest-loser pricing” (HLP).<sup>21</sup> The final rules—HLP on both  $A$  and  $B$ —correspond to the rules originally proposed in Klemperer (2008).<sup>22</sup> We call these rules PMA-HLP.

Bids strictly above the clearing prices are fully allocated, and bids on the margin for a good are rationed pro-rata.

Table 2 also shows the starting value of the relative supply curve as it varies over the period, separately for the 3-month and 6-month term auctions. The relative supply curve is an increasing step function, with step length equal to 0.1% of the maximum supply.

In Table 2, the rule changes are shaded in grey.<sup>23</sup>

<sup>21</sup>Under LWP, the price for each good is equal to the lowest winning bid on that good. Under HLP, the price for good  $A$  is equal to the highest losing bid on that good, or 0bps if all bids on good  $A$  are fully allocated. The price for good  $B$  under HLP equals the clearing price for good  $A$  plus the value of the relative supply curve at the quantity allocated of good  $B$ .

Using LWP introduced slight welfare losses in the earlier period under specific circumstances. When the auction was not fully allocated but the clearing price of good  $A$  was strictly positive, the implied minimum price that the BoE was willing to accept on the margin for good  $B$  was raised above the true price that the BoE was willing to accept. In four auctions, this caused the quantity allocated of good  $B$  to be slightly below the optimal quantity, so welfare was not maximal.

To evaluate the significance of this change, we repeat the surplus difference calculations using HLP for the PMA price determination in all auctions. The results are very similar to the corresponding results for the actual PMA, shown in Appendix A.4. The net welfare differences are almost identical and the transfers from bidders to the BoE under the PMA are slightly smaller, as the clearing prices are on average lower when using HLP.

<sup>22</sup>HLP is preferable to LWP as small bidders, at least, then have no incentive to misrepresent their preferences. If a bidder expects LWP, they have an incentive to underreport their true demand, due to the possibility of their bid determining the market clearing price. Regardless of the size of their demand, a bidder’s incentive to distort their bid away from their valuation is larger under LWP than HLP.

<sup>23</sup>In many cases, the changes in price determination rules are irrelevant. However, in cases in which the auction is uncovered, the highest losing bid on good  $A$  is interpreted as equalling 0 (i.e. the reserve price), which may be lower than the lowest winning bid. Similarly, the price that the BoE was willing to accept on the margin for good  $B$  may be lower than the lowest winning bid.

**Usage** Usage of the auctions substantially decreased from June 2012.<sup>24</sup> The greater demand for funds in the earlier auctions increases the scope for misallocation and welfare losses if the operation is not appropriately designed. The average differences in surpluses between designs are weighted by the quantity allocated in the PMA.

### 3 Method

#### 3.1 Truthful bidding

We assume that bidders submit bids that describe their true willingness to pay for the two goods. If bidders believe themselves to be price takers, so that their bids do not affect the market clearing prices, bidding truthfully is an optimal strategy.

Our assumption is motivated by the observation that a large proportion of bidders submit a single bid, constituting a flat demand curve. If bidders expect to be pivotal and act strategically, bidders optimally submit multiple staggered bids (under reasonable assumptions on bidders' preferences), and therefore engage in "laddering".<sup>25</sup> Further, the number of potential participants is known to be large,<sup>26</sup> and neither the actual number of participants nor the BoE's relative supply curve for any auction is ever disclosed to the bidders,<sup>27</sup> making optimal strategic behaviour difficult to determine.<sup>28</sup> Grace (in prep., a) finds that truthful

---

<sup>24</sup>The average cover in June 2012 – January 2014 is 0.4% for good *A* and 3% for good *B*, compared to 74% and 23% respectively in June 2010 – May 2012. A plausible reason for this is an increase in the availability of funding from other sources. The Funding for Lending Scheme was introduced by the BoE in July 2012, allowing borrowing of up to four years. In addition, the BoE's Extended Collateral Term Repo (ECTR) operation, consisting of monthly auctions for 6-month loans is another source, although this was only activated once (in June 2012). Increased availability of funds from the euro area also plausibly reduced demand in the ILTR auctions.

<sup>25</sup>A bidder with market power has an incentive to misreport their preferences. By understating their willingness to pay, they push down the expected market clearing price and reduce their expected total payment for the allocation of funds. A fully strategic bidder trades off this gain with the possibility of also losing in cases in which they would prefer to win.

<sup>26</sup>Participants of the BoE's Sterling Monetary Framework with access to the BoE's Open Market Operations (OMO) are eligible to register to participate in the ILTR auctions. The number of institutions with OMO access (i.e. the number of potential participants) was fairly stable from October 2009 (before the ILTRs were introduced) to January 2014 and ranged from 48 to 52.

<sup>27</sup>Bidders could have inferred some features of the supply curve from the information publicly released after the auctions. For example, in February 2011, the 6-month term auction is uncovered but some bids on good *B* are unallocated and the marginal bid on good *B* is rationed. Given this, the bidders could infer a point on the 6-month term relative supply curve: at the total quantity allocated of good *B* (33.3%), the BoE's relative supply equals the difference in the clearing prices for *B* and *A* (52bps). However, the amount of information revealed over time was limited and the supply curves shifted over time as shown in Table 1, limiting the scope to act on this information strategically.

Bidders could also have inferred the number of participants to be small in some of the later auctions (from October 2012 onwards), based on the publicly revealed information.

<sup>28</sup>Because bidders for both goods bid within the same auction in the PMA, the PMA increases competition

bidding appears to be a reasonable approximation in the ILTR auctions.

Under this assumption of truthful bidding, the PMA maximises welfare.

## 3.2 Alternative auction rules

We use the PMA as a welfare-maximising benchmark to compare two alternative auction rules.

### 3.2.1 Comparison 1: Separate simultaneous auctions

**Separate simultaneous auctions with fixed quantities supplied (SSA)** We first compare the PMA to running two separate auctions simultaneously, one for each good, with the quantity supplied of each good, A and B, announced and fixed prior to the auction.

For each of the 3-month and 6-month term auctions, we fix the total supply available across both goods to the maximum supplies of £5 billion and £2.5 billion, respectively. We then vary the fixed percentage of this total that is supplied of good  $B$  from 0% to 100%. Under each of these auction rules, we calculate the market equilibrium with reserve prices equal to the starting point of the supply curve in each period, shown in Table 1, assuming that bidders' bidding behaviour remains unchanged. The market clearing prices equal the highest losing bids for the two goods. We then compare the realised surpluses and net welfare relative to the welfare-maximising PMA benchmark.

The SSA constrains the BoE relative to the PMA, eliminating the flexibility in supply that the PMA allows both through the joint determination of clearing prices and through the relative supply curve.

Importantly, the SSA allocation does not depend on the relative supply curve, and so there may be outcomes in which the BoE is supplying good  $B$  at a price below what it would be willing to accept on the margin. Large welfare losses can therefore arise in SSAs in which the fixed quantity allocated of good  $B$  is large, for auctions with high demand for good  $B$ .

**Reserve prices** To test for robustness of the results, we also consider variants of the SSA, in which there are (i) no reserve prices, (ii) reserve prices of 5bps for good  $B$  across the whole period and (iii) reserve prices of 5bps for all 3-month auctions and 15bps for all 6-month auctions relative to the SSA, which may reduce the incentive to bid strategically.

auctions. The SSA that we focus on is the most favourable one in terms of both welfare and BoE surplus. In this design, the reserve prices change over time in line with the observed changes in the actual supply curve, despite the BoE making these changes in response to signals from the PMA that would have been unobserved in an SSA. The welfare losses from this SSA can therefore be thought of as a lower bound on the welfare losses resulting from an SSA.

We also discuss the impact of running the two separate auctions sequentially.

### 3.2.2 Comparison 2: Reference price auction

**Reference price auction (RPA)** Second, we compare the PMA to running a single-round auction in which a pair of notional prices for goods  $A$  and  $B$  are fixed in advance, and the highest bids across the goods relative to their reference prices are accepted, with the total supply available across both goods fixed and announced prior to the auction. This design has been used in financial markets, for example, the BoE and Federal Reserve have used RPAs for their quantitative easing purchase schemes (Song and Zhu, 2018; Bank of England, 2022).<sup>29</sup> The notional prices might reasonably be fixed at “best guesses” of the fair market prices. For example, the BoE’s gilt purchase operations and Corporate Bond Purchase Scheme are RPAs in which the reference price for each bond is its secondary market mid yield at the close of the auction (Bank of England, 2022).

To implement an RPA, we normalise the reference price of good  $A$  to zero, subtract the reference price of  $B$  from the bids on good  $B$ , and find the market clearing allocation as if all bids are for a single good. This then implies clearing prices for goods  $A$  and  $B$ , equal to the highest losing bids for the two goods.<sup>30</sup> The fixed reference prices mean that the difference in clearing prices for goods  $A$  and  $B$  equal the difference in reference prices. The marginal winning bid is restricted to be non-negative. This effectively imposes a reserve price for good  $A$  of 0bps and a reserve price for good  $B$  equal to its reference price.

Given the nature of the two “goods”—loans of funds secured by Levels A and B collateral—this can be interpreted as imposing a fixed spread between loans secured by collateral assets with different liquidity levels.

---

<sup>29</sup>Armantier, Holt and Plott (2013) report that the US Treasury also considered using an RPA to purchase “toxic assets” (namely mortgage-backed securities) in the \$700 billion Troubled Asset Relief Program (TARP) in 2008.

<sup>30</sup>Bids strictly above the respective clearing prices are fully allocated and bids on the margin in a fully allocated auction are rationed pro-rata with a common rationing coefficient across goods  $A$  and  $B$ .

We also compare the PMA to posting a pair of fixed prices, with the total supply available fixed and bidders with bids above the posted prices rationed pro-rata. This can be viewed as an RPA in which the auction prices are also fixed.

### 3.3 Surpluses, and welfare

We calculate the auction outcomes under the different rules and calculate the difference in total net bidder surplus (denoted bidder surplus), BoE surplus, BoE revenue, and net welfare between the actual PMA and the alternative auction.

**Bidder surplus** A bidder’s bids for a particular good can be aggregated to form their individual demand curve for that good, which is equal to their marginal valuation curve under the assumption of truthful bidding. For each bidder, net bidder surplus therefore equals the sums of the areas under their demand curves for goods  $A$  and  $B$  up to the quantities allocated, minus the total payment that they make. Bidder surplus, equal to total net bidder surplus, is the sum of individual net bidder surpluses.

**BoE surplus** The maximum supply in each auction and the position of the relative supply curve are “pinned down by [the BoE’s] preferences — namely to provide liquidity insurance to the banking system at a price that doesn’t undermine the incentive to prudently manage liquidity” (Fisher, 2011, p.12). We interpret this to mean that the supply curves reflect the minimum prices that the BoE is willing to accept for the loans, as functions of the allocation. In particular, for any total quantity less than the maximum supply, the price that the BoE is willing to accept for good  $A$  is zero and the price that it is willing to accept to supply a marginal unit of good  $B$  equals the height of the relative supply curve. We also assume that, for each auction, the price at which the BoE is willing to supply any quantity in excess of the maximum supply is infinite, so the total quantity allocated can never exceed the maximum supply.<sup>31</sup> For good  $A$ , BoE surplus is therefore the market clearing price for good  $A$  multiplied by the quantity allocated of good  $A$ . For good  $B$ , it is the difference between the market clearing price for good  $B$  and the BoE’s relative supply curve, up to the quantity allocated of good  $B$ . See Appendix A.3 for further discussion.

**BoE revenue** BoE revenue is measured by the sum of the quantity allocated of good  $B$  multiplied by the market clearing price for good  $B$  plus the quantity allocated of good  $A$  multiplied by the market clearing price for good  $A$ .


---

<sup>31</sup>The relative supply function is a step function, over a fine quantity grid, with steps at each 0.1 percentage point of the maximum supply. For quantities allocated between the two steps, the price that the BoE is willing to accept equals the lower bound of the two steps.

**Net welfare** Total net welfare equals the sum of bidder surplus and BoE surplus.<sup>32,33</sup>

Figure 2 illustrates these surpluses for each good in the PMA, at market clearing prices  $(p_A, p_B)$  and quantities  $(q_A, q_B)$ . Surpluses in the alternative designs are analogous.

Figure 2: Surpluses at the PMA market clearing prices,  $(p_A, p_B)$ , and quantities,  $(q_A, q_B)$



#### 4 The SSA

We first compare the surpluses resulting from the actual PMA and the SSA, as the proportion of total quantity supplied of good  $B$  varies from 0% to 100%.

Figure 3 show the differences in auction surpluses between the actual PMA and SSA, averaged over the 3-month and 6-month term auctions in June 2010 – January 2014, weighted

<sup>32</sup>The welfare-maximising total quantity allocated may be less than the maximum supply even if the auction is “covered” (i.e. total demand across both goods at least equals the maximum supply). The reserve price for good  $A$  is 0bps, and the relative supply curve constrains the welfare-maximising allocation of good  $B$  so that the difference between the lowest winning bids for goods  $B$  and  $A$  weakly exceeds the difference in prices that the BoE is willing to accept on the margin for goods  $B$  and  $A$  (equal to the height of the relative supply curve).

<sup>33</sup>To illustrate, consider the July 2010 auction. In this auction, the clearing prices for goods  $A$  and  $B$  are 1bp and 26bps, respectively. 17.2% of the maximum supply is allocated of good  $B$ , and 82.8% allocated of good  $A$ . The marginal bid on good  $A$  is 1bp, yielding marginal net welfare of 1bp. The marginal bid on good  $B$  is 26bps and the price that the BoE is willing to accept on the margin for good  $B$  (equal to the relative supply curve evaluated at 17.2%) is 25bps. So the marginal net benefit of allocating a unit to good  $B$  is also 1bp. Equating the two marginal net benefits maximises net welfare.

by the quantities allocated in the PMA. The surpluses in an auction are measured in basis points of the quantity allocated by the PMA in that auction.

There are five key features of the results.

**Feature 1 Welfare maximisation could almost have been achieved in the SSA given observed demand**


In Figure 3, the average net welfare difference between the SSA and PMA is U-shaped, strictly decreasing to approximately zero and then strictly increasing. This shows the welfare loss of imposing the SSA constraint on the auction. If the quantities supplied of good  $B$  were fixed at 17.5% for 3-month terms and 25% for 6-month terms, the average welfare loss from running the simpler SSA would have been essentially zero.

This suggests that, had the BoE been able to forecast these optimal quantities, implementing the SSA would have almost maximised welfare over the period. However, this seeming strength of the SSA is limited for two reasons:


1. **Forecasting** These optimal fixed quantities were very uncertain ex ante (Fisher, 2011, p.8), so the BoE could not have implemented the best-performing SSAs in practice and the welfare losses of forecast error can be significant (see Feature 2). Moreover, if the BoE had adjusted the fixed quantities based on the outcomes of the initial auctions, the welfare losses might have been small over the sample period because the optimal fixed quantities turned out to be fairly constant. However, it could not have relied on this stability ex ante. (This uncertainty in both the optimal quantities and their stability was one motivation for the introduction of the ILTR auctions, replacing the LTR operations, in 2010.)
2. **Endogenous response** The ILTR auctions were introduced and designed to flexibly and systematically respond to financial stress. The composition of demand for funds in the June 2010 – January 2014 period was relatively stable, so that the optimal quantity supplied of good  $B$  was relatively stable across auctions. In a time of stress, demand for good  $B$  is likely to rise relative to good  $A$ , so that the optimal quantity allocated of  $B$  would likely rise. If unanticipated, an SSA could generate significant welfare losses relative to the PMA, and if anticipated, a publicly announced change in the composition of the SSA could add to market instability.

Figure 3: Difference in auction surpluses, PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term



**Feature 2 But the welfare losses could have been large if the BoE had misjudged the correct proportion to assign to good  $B$**

Figure 3 shows that the average net welfare loss from the SSA can be significant. For example, if the BoE had fixed the quantity supplied of good  $B$  to be 50% of the maximum supply, this would have led to an average welfare loss of over 2bps per auction for both 3- and 6-month terms, equal to 32% of the average welfare generated by the PMA: if the BoE had fixed the quantity supplied of good  $B$  to be 2/3 of the maximum supply, the average welfare losses would have been approximately 2.4bps and 4.6bps (41% and 47% of PMA welfare) respectively for the 3- and 6-month term auctions.<sup>34</sup>

**Feature 3 Net welfare changes smoothly but changes in bidder surplus and BoE surplus and revenue are lumpy**

The average difference in net welfare changes smoothly as the fixed quantity supplied of good  $B$  rises. In contrast, the average revenue, bidder surplus and BoE surplus differences are lumpy.

The allocation changes approximately continuously<sup>35</sup> as we change the proportion allocated to good  $B$ , but the clearing prices are discrete and so change discontinuously. In addition, at the optimum of 17.5% for 3-month terms and 25% for 6-month terms, welfare losses are second order whereas the transfers between agents are first order. Net welfare is determined by the allocation, whereas bidder surplus, BoE surplus and revenue are determined by both the allocation and clearing prices.

**Feature 4 Large transfer from BoE to bidders in some SSAs relative to the PMA**

In Figure 3, the average BoE surplus is uniformly higher in the PMA than the SSA across all fixed quantities supplied of good  $B$ . The average bidder surplus is higher in the PMA than the SSA for low fixed quantities supplied of good  $B$  (up to 7.5% fixed quantity in the

---

<sup>34</sup>Observe also that if none of the maximum supply were allocated of good  $B$ , the average welfare losses would have been 3.1bps and 8bps (54% and 66% of PMA welfare) respectively for the 3- and 6-month term auctions. This gives some indication of the foregone welfare had the BoE not extended the acceptable range of collateral to include Level B collateral in the liquidity operations from June 2010. (This estimate is an upper bound), because some bidders who bid against Level B collateral in the PMA might perhaps have submitted bids against Level A, had Level B collateral not been accepted. However, the rare use of paired bids suggests this substitution behaviour might have been limited.)

<sup>35</sup>The smallest unit of allocation is £100,000, equal to 0.002% and 0.004% of the maximum supply in the 3-month and 6-month term auctions respectively, so the allocation can be treated as approximately continuous.

3-month auctions and 20% in the 6-month auctions). It is uniformly higher in the SSA than the PMA across all fixed quantities higher than these respective thresholds. In particular, average bidder surplus is 1.5bps higher in the best-performing SSA than the PMA for the 3-month auctions and marginally higher for the 6-month auctions. This is the case despite the fact that the efficiency loss of the best-performing SSA is small. The allocation chosen by the PMA is not constant, and small differences in the allocation between the PMA and best-performing SSA can create large changes in bidder and BoE surpluses through changes in the clearing prices.

There are two reasons for these differences in bidder surplus and BoE surplus:

1. Misallocation of goods  $A$  and  $B$
2. Transfer of surplus between the BoE and bidders through differences in the market clearing prices

First consider the SSAs with very low fixed quantities supplied of good  $B$  and corresponding high fixed quantities supplied of good  $A$ . In these cases, good  $A$  is over-allocated and good  $B$  is under-allocated relative to the optimal quantities. There is large misallocation in the SSA, causing both bidder surplus and BoE surplus to be higher in the PMA. The high fixed quantity allocated of good  $A$  also pushes down the clearing price for good  $A$ , yielding a transfer from the BoE to bidders in the SSA.

For higher fixed quantities supplied of good  $B$ , bidder surplus is lower and BoE surplus is higher in the PMA relative to the SSA. As the fixed quantity allocated of good  $B$  rises beyond a certain point, good  $B$  is over-allocated and  $A$  is under-allocated, holding the total quantity fixed. In addition, the SSA allocation is determined independently of the BoE's relative supply curve so that there may be units allocated of good  $B$  for which the price that the BoE would be willing to accept exceeds the marginal benefit of allocating the loans. This generates further loss in BoE surplus and gain in bidder surplus in the SSA relative to the PMA.

The transfer of surplus between the BoE and bidders in the SSA may also be substantial at these higher fixed quantities supplied of good  $B$ . The market clearing prices in the SSA are determined entirely by the highest-losing bids, independent of the prices that the BoE would be willing to accept for the marginal loan. As the fixed quantity supplied of good  $B$  rises, eventually all of the bids on good  $B$  are fulfilled, so that the clearing price for good  $B$  falls to the reserve price. In contrast, the market clearing price in the actual PMA either

equals the lowest winning bid or the maximum of the highest losing bid and the price the BoE is willing to accept for the marginal loan. The market clearing price for good  $B$  in the SSA is typically far below that in the actual PMA, generating a large transfer from the BoE to the bidders.

In addition, the transfer from BoE to bidders is partly driven by the price determination rules. Recall the PMA moves from an LWP rule to an HLP rule (which is a function of the price that the BoE is willing to accept on the margin). In contrast, the SSA adopts an HLP rule throughout the entire period, which is independent of the supply curve. By construction, the LWP increases BoE surplus for a given allocation (assuming truthful bidding). To investigate this, we also repeat the welfare calculations for a PMA with HLP for the entire period, where the highest loser in the PMA case is the maximum of the price that the BoE is willing to accept on the margin and the highest losing bid. The differences in BoE surplus and bidder surplus are slightly attenuated, especially around the range of minimum welfare loss.<sup>36</sup>

#### **Feature 5 Reserve prices do little to reduce misallocation**

One key source of welfare loss in the SSA is an allocation of good  $B$  in excess of the optimal quantity so that the price that the BoE would be willing to accept on the margin exceeds the marginal bid (equal to the marginal bidder valuation).

A crude way to mitigate this welfare loss is to set reserve prices in the SSA. We assumed the SSA has reserve prices equal to the minimum spreads of the relative supply curve in each auction, shown in Table 1. We also calculate the welfare losses under three different reserve prices: (i) no reserve prices, (ii) reserve prices of 5bps for good  $B$  across the whole period, and (iii) reserve prices for good  $B$  of 5bps for all 3-month auctions and 15bps for all 6-month auctions. The average surpluses for 3- and 6-month auctions are very similar in all four cases.

#### **4.1 Separate auctions run sequentially**

Running separate auctions sequentially would almost certainly yield similar results to running them simultaneously, because the fact that paired bids were hardly used suggests that bidders would bid the same way in each auction whether they were sequential or simultane-

---

<sup>36</sup>The benefit to bidders of the PMA is small (but positive) when comparing the PMA-HLP to an SSA with a small quantity supplied against Level B, but the bidders receive essentially no benefit from the PMA when comparing with the actual version of the PMA used (with LWP in the earlier auctions, rather than HLP).

ous.<sup>37,38</sup>

## 5 The RPA

Second, we compare the surpluses between the PMA and the reference price auction (RPA) as the reference price of B varies from 0bps to 60bps. The RPA seems a particularly natural comparator since the BoE uses the design for its quantitative easing purchases and used a somewhat related design in the LTR auctions prior to the implementation of the ILTR auctions in June 2010.<sup>39</sup>

Figure 4 show the differences in auction surpluses between the PMA and RPA, averaged over the 3-month and 6-month term auctions in June 2010 – January 2014, weighted by the quantities allocated in the PMA. As for the SSA comparison, the surpluses in an auction are measured in basis points per loan of the quantity allocated by the PMA in that auction.

There are five key features of the results.

---


<sup>37</sup>Paired bids allowed bidders to express one-for-one substitute preferences, but bidders with more complex substitute preferences could still have gained by using them, so the near-absence of paired bids suggest that bidders viewed the goods independently. Moreover, either one-for-one substitutability or non-substitutability between the goods appears most plausible: they are measured in the same units and generate the same amounts of liquidity. Indeed, the BoE designed the ILTR auction anticipating bidders to express one-for-one substitutability, if at all.

<sup>38</sup>Running the auctions sequentially permits some adjustment between auctions by either bidders or the auctioneer. The information bidders learn in the first auction could affect behaviour in the second auction, but the effect seems likely to be small. The auctioneer could somewhat improve the performance of the sequential auctions by making any unused supply from the first auction available in the second auction.


<sup>39</sup>In the LTR operations, loans over a 3- to 12-month term were auctioned against both Level A and B collateral sets. This was a discriminatory auction in which bids were accepted in decreasing order, with a reserve price of 50bps applied exclusively to bids for loans secured by Level B collateral.

Figure 4: Difference in auction surpluses, PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term



## **Feature 1 Welfare could not have been maximised with RPAs given observed quantities**

The average net welfare difference between the actual PMA and the RPA does not fall to zero in either the 3-month or 6-month auctions over the period studied. Figure 4 shows that, for the entire range of reference prices of  $B$ , the average welfare loss of the RPA is at least 0.6bps and 0.8bps, respectively for the 3- and 6-month auctions. This is equivalent to 10% and 14% of the average welfare generated by the PMA, respectively.

This results from (i) aggregation and (ii) persistent welfare losses at the individual auction level.

**Aggregation** Firstly, the welfare losses in the RPAs do not all reach their minima at one particular reference price of  $B$ , unlike the losses in the SSA which all reached their minima at one particular fixed quantity. The optimal price difference between goods  $A$  and  $B$  varies across auctions so that the BoE could not commit to a reference price of  $B$  that would achieve the welfare-maximising allocation over time. This is the case despite the environment being relatively stable over time.

**Auction-level persistent welfare losses** Moreover, there are some auctions in which the welfare loss never reaches zero even for the individual auction. This is a more general issue with the RPA, considered in the next section.

## **Feature 2 The RPA rarely maximises welfare and usually creates significant losses**

The welfare losses of the RPAs observed in the period studied is a general feature of the RPA. The reason is that bidders typically submit single bids for large quantities, rather than a set of smaller bids at laddered prices, so even a small change in the reference price of  $B$  can cause a large change in the allocation, from a quantity of good  $B$  which is much less than the optimal one, to a quantity of good  $B$  which is much more than is optimal. Thus even for a given auction, an RPA with a reference price of  $B$  close to the optimal price difference for that auction can still generate large welfare losses.<sup>40</sup> Since the optimal difference between

---

<sup>40</sup>Indeed, even if the fixed price is optimal, there will still in general be welfare losses due to suboptimal rationing of marginal bids. For example, consider the January 2011 auction: even an RPA which turned out, after the fact, to be the best possible one would have generated a welfare loss equivalent to over 2bps of the total quantity allocated that month by the PMA relative to the welfare-maximising outcome. The optimal reference price of  $B$  would have been 26bps, but the RPA then accepts all the marginal bids on good  $B$ , which represent 29% of the maximum supply (as the auction is uncovered, so the common RPA rationing coefficient equals 1). By contrast, the optimal allocation accepts just 17.2 percentage points of

the clearing prices for goods  $A$  and  $B$  varies across auctions (this was part of the point of introducing the PMA), any RPA with a price difference that is constant across auctions will lead to significant losses, even when the environment is relatively, but not perfectly, stable.<sup>41</sup>

**Feature 3 Welfare losses, BoE revenue and BoE surplus change discontinuously across reference prices but bidder surplus changes smoothly**

Figure 4 shows that the differences in net welfare, and BoE revenue and surplus of the RPAs change discontinuously as the reference price of  $B$  changes, yet the difference in bidder surplus smoothly increases.

**Welfare losses** The lumpiness of the welfare loss is driven by the same property of the RPA that causes the mechanism to be suboptimal in general. As the reference price of  $B$  changes, the ordering of the bids relative to their reference prices changes. In particular, the ordering of bids on goods  $A$  and  $B$  changes. If the ordering changes so that a large bid on good  $B$  shifts from being allocated on the margin to fully unallocated, the overall quantity allocated of good  $B$  decreases significantly, causing a large discrete change in welfare.

**BoE surplus and revenue** Since the allocation between goods  $A$  and  $B$  changes in jumps, BoE revenue and BoE surplus also change discontinuously as the reference price of  $B$  changes.

**Bidder surplus** Bidder surplus at the margin is equal to zero for bidders on both goods  $A$  and  $B$ , because, for each good, the marginal bidder's own bid sets the clearing price. So, a small change in the allocation between goods  $A$  and  $B$  that does not change the prices does not cause large shifts in bidder surplus. Moreover, the price of  $A$  changes very little as the reference price of  $B$  changes and often equals zero, so the price for  $B$  increases, and bidder

---

these marginal bids on good  $B$  at the spread of 26bps.

For a *given* auction, a fixed price together with a sophisticated rationing rule that mimicked the behaviour of the PMA would, of course, resolve the problem for that auction. However, this would not help much across multiple auctions where different fixed prices would be needed, and where even a slightly “wrong” fixed price for an auction could lead to substantial welfare losses.

<sup>41</sup>Note that this contrasts with the SSA which can perform relatively well in a fairly stable environment if it is correctly calibrated. However, this gives no reason to suppose that either the RPA or SSA would perform better than the other in an environment with significant shocks. (With larger shocks, their relative performance would depend upon the issues discussed in Weitzman (1974) (with price-difference in our context taking the role of price in that paper).)

surplus decreases, almost smoothly with B’s reference price.<sup>42,43</sup>

In sum, the impacts of small changes in the RPA reference price are opposite to the results of the SSA, described in Section 4. In the SSA, small changes in the quantities of loans supplied against the two collateral sets cause small changes in the allocation, but can cause large changes in the clearing prices, and hence bidder surplus; in the RPA, small changes in the price difference between the two collateral sets can cause large changes in the allocation, but cause only small changes in the clearing prices and bidder surplus.

#### **Feature 4 Large transfer from BoE to bidders in some RPAs relative to the PMA**

In Figure 4, average BoE surplus and revenue are almost uniformly higher in the PMA than the RPA for all reference prices, but bidders as a group only share the welfare gain of the PMA relative to the RPA for high reference prices of B.

When the reference price of B is very small, the market clearing price for good *B* in the RPA is typically far below the PMA clearing price for B. More bidders win good *B* in the RPA and, holding the quantity sold constant, bidders benefit more from the RPA than the PMA so that the RPA increases bidder surplus. Conversely, the BoE benefits from the higher revenue in the PMA, and the PMA only allocates bids at prices that the BoE is willing to accept (as measured by the supply curves), whereas the RPA overallocates B relative to the PMA’s welfare-maximising quantity, at prices below the supply curve. Because of these two factors, BoE surplus in the PMA is consistently higher than in the RPA.<sup>44</sup>

At the other extreme, when the reference price of B is very large, the difference between bid

---

<sup>42</sup>Figure 4 illustrates this. When the reference price of good *B* is zero, bidder surplus is much higher in the RPA than in the PMA. In this case, bids on good *A* and good *B* are treated interchangeably and, because the bids on B are almost always higher than bids on A, they are allocated first. Since the cover ratio of good *B* is always less than 1, all bids on good *B* are allocated in all auctions. In almost all auctions, the marginal bid is on good *A*, so that the clearing prices for both goods *A* and *B* is around 0-2bps (far below the typical PMA clearing price for good *B*). This therefore generates a huge bidder surplus in the RPA relative to the PMA. As the reference price of B increases, the ranking of the bids in each RPA changes very little so that the allocation changes very little. The clearing price of B therefore increases almost smoothly with the reference price, causing bidder surplus to smoothly decrease in the RPA.

<sup>43</sup>As the reference price of B increases above 26bps for the 3-month auctions and above 46bps for the 6-month auctions, the bidder surplus in the PMA exceeds bidder surplus in the RPA. This follows from the fact that the marginal winning bid is constrained to be non-negative. When the reference price of B is very large, the bids on good *B* are below their reference price and therefore go unallocated in the RPA. This leads to a loss of bidder surplus relative to the PMA, in which some bids on good *B* are typically allocated.

<sup>44</sup>This is analogous to the result in Section 4 that BoE surplus is uniformly higher in the PMA than the SSA. In that case, the difference acts through the price determination rule — the PMA clearing prices depend on the relative supply curve, whereas the SSA clearing prices do not.

and reference price for many bids on good  $B$  is negative, so that they are unallocated. The RPA therefore underallocates  $B$  relative to the PMA, which allocates the efficient quantities, and raises the price paid for  $B$ , so that bidder surplus is lower and BoE surplus is higher in the RPA than in the PMA.

As in Section 4, this transfer from BoE to bidders is partly driven by the price determination rules, as the PMA moves from an LWP rule to an HLP rule. However, we calculated the average differences in surpluses for the PMA with an HLP rule for the entire period and the RPA and found only a small change in transfers.<sup>45</sup>

### **Feature 5 Fixed prices could also cause large welfare losses**

A simpler alternative mechanism would be for the BoE to fix two prices, one for each of goods  $A$  and  $B$ , at which bidders could borrow funds. However, this would likely cause significant welfare losses, in excess of those estimated for the RPA over this period for two reasons. Most importantly and akin to the difficulty of choosing the reference price of  $B$  in the RPA, the BoE would have no ability to predict the optimal fixed prices prior to the auction, so that the posted prices would possibly be far from optimal. Secondly, if the total quantity of funds was capped and excess demand was rationed, the loans would not go to the bidders that value it most, causing welfare losses.<sup>46</sup>

To illustrate this, we compared the surpluses for the PMA to posting a pair of fixed prices, at a range of different levels, with a fixed supply, set equal to the BoE's maximum supply, and winning bids rationed pro-rata. Figures B.1 and B.2 show these differences in surpluses as the fixed price for Level  $B$  changes, holding constant the fixed price of Level  $A$ , averaged over the 3-month and 6-month term auctions, respectively. These are qualitatively similar to the differences in surpluses between the PMA and RPA. We would expect the differences between the PMA and fixed prices to be quantitatively larger in a period of greater stress.

---

<sup>45</sup>There are some minor differences. For example, the benefit to bidders of the PMA is small (but positive) when comparing an RPA with a large reference price of  $B$  to the PMA-HLP, but is zero when comparing with the actual version of the PMA used.

<sup>46</sup>If, instead, the total supply of funds was uncapped, posting fixed prices may lead to an extremely adverse welfare outcome, depending on the prices that the BoE would be willing to accept to supply funds beyond the BoE's stated maximum supply. For example, if the prices that the BoE was willing to accept for goods  $A$  and  $B$  were constant beyond the stated maximum supply and equal to the prices at that maximum supply, there might be a large adverse welfare outcome, especially if the BoE underestimates the demand for Level  $B$  at the posted price.

## 6 Welfare and surpluses in periods of stress

Demand was relatively stable in the period studied, so that the welfare benefits, arising from the flexibility of the allocation in the PMA, were fairly muted. To investigate the PMA's relative performance in a period of greater stress, we compare the surpluses and welfare of the PMA with the SSA and RPA given simulated bids. We transform the observed bids by multiplying the observed quantities demanded by 4 and 8 for the bids on Levels A and B, respectively, and multiplying the observed bid prices submitted for Level B by 4. This is intended to loosely reflect the pattern of demand at the beginning of the COVID-19 pandemic in March 2020. Figures B.3 and B.4 repeat the calculations of the differences in surpluses using these simulated bids.

The welfare gain of the PMA is qualitatively similar, but quantitatively significantly larger in absolute terms, in this exercise simulating stress, reflecting the scaling up of benefits arising from the flexibility of its allocation across collateral types.

As in the main analysis, the BoE's surplus is always larger—and is now very much larger—in the PMA than in either the SSA or the RPA in this experiment. Importantly, because the PMA's allocations respond to the BoE's supply preferences, the BoE obtains (weakly) positive surplus from every unit it allocates, which is not always the case in the SSA and RPA, whose allocations are insensitive to its preferences.

However, the impact on bidders' surplus differs across comparators: the bidders gain from the PMA over a larger range of SSAs, but over a smaller range of RPAs.

Figure B.3 shows that the PMA now gives bidders more surplus (in aggregate) than all SSAs that allocate less than 45% to Level B for 3-month terms, and all that allocate less than 75% for 6-month terms. For the original demands, the PMA only gave bidders more surplus than SSAs that allocated less than 10% and 20% to Level B, respectively. The reason for these results is that the relatively flat demand for loans secured by Level A (which is preserved in the transformations) means bidders can never gain much surplus on Level A. So when there are a large number of high bids on Level B (as would be expected during a period of stress), bidders gain more surplus from the PMA, which will allocate a large share of the maximum supply to B in these circumstances, than from an SSA, unless the SSA also allocates a large share to Level B.

By contrast, Figure B.4 shows that the PMA now only gives bidders more surplus than RPAs with a reference price of B higher than 78bps for 3-month terms and higher than 68bps for

the 6-month terms. For the original demands, the PMA gave bidders more surplus than RPAs with a reference price of B higher than 26bps and 48bps, respectively. The reason is that in periods of greater stress bids on B increase relative to A, so for any reference price of B, the RPA accepts more on Level B, generating larger net surplus to the winning bidders. The PMA also accepts some more on Level B relative to the original demands, because it is more efficient to do so when bids on B are relatively high. But the change in the RPA is typically larger—to the benefit of the bidders—as the PMA also takes the BoE’s supply preferences into account. So, under stressed conditions, bidders prefer a PMA to an RPA only if the RPA has a very high reference price of B.

We experimented with various alternative transformations of observed demand and they yielded broadly similar results. However, the results about the bidders’ preferences between the designs may not generalise beyond the specific case of two eligible collaterals of which one is Level A, because these results seem to be caused by the flat demand on Level A.

## **7 BoE-surplus-maximising and BoE-revenue-maximising PMA**

The PMA used by the BoE in the period studied is designed to maximise total welfare. We find that changing this objective to maximising BoE surplus has limited qualitative impact on the results. Figures B.5 and B.6 compare the surpluses of the SSA and RPA to a PMA, in which the allocation adjusts to maximise BoE surplus.

The welfare gain of the PMA is broadly similar both qualitatively and quantitatively to the case in which the PMA maximises total welfare. This reflects the importance of the PMA’s responsiveness to the BoE’s supply preferences in driving the welfare gains relative to the SSA and RPA.

The distributional impact, that is, the impact on the bidders’ and BoE’s respective surpluses, is qualitatively similar but quantitatively larger to the main analysis. This is unsurprising as the BoE-surplus-maximising PMA systematically favours the BoE relative to the bidders, amplifying the observed benefits to the BoE of the PMA over the sample period.

The comparisons for a PMA that maximises BoE revenue are broadly similar. (The welfare gain of a PMA which maximises BoE revenue, i.e. the total payment from bidders to the BoE, is slightly smaller than that of the BoE-surplus-maximising PMA, whereas the transfer of surplus from bidders to the BoE is also typically slightly smaller.)

## 8 Broader benefits of the PMA

Unlike the SSA and RPA, the PMA adjusts both the price and the quantity allocated of each good in response to the bidders' and BoE's preferences (expressed by the bids and the supply curves, respectively). The benefits of this automatic adjustment can be split into the direct welfare gains of meeting the liquidity needs of individual participants, and the externalities that Fisher (2011) describes of providing liquidity insurance to the system as a whole, avoiding moral hazard, and informing the BoE about stress in the market (in addition to communicating the "correct" price to the market).<sup>47</sup>

One measure of the PMA's performance in terms of the overall benefits would be how far the outcome of an alternative design would be from the optimal prices and quantities. To the extent that the BoE's supply curves express the optimal price-quantity relationship, the PMA is designed to allocate the "correct" quantity at the "correct" price but the alternatives do not necessarily do so.<sup>48</sup> The alternatives' allocations would not just be inefficient as measured by the spreads on the loans, but also for society more broadly. How one measures this is a matter of choice — it could be the price difference from the optimum, the quantity difference, or the product of the two. Our measure is roughly proportional to the product of the two.<sup>49</sup>

The difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept on the loans allocated in the auctions does not capture all the benefits of the PMA. Our measure is therefore a conservative estimate of the overall benefits.

---

<sup>47</sup>Only some of these broader benefits are relevant to the comparison of the PMA to each alternative. The benefit of providing system-wide liquidity insurance is particularly important in the comparison of the PMA to an alternative which supplies too little funds against Level B, but less important when comparing the optimum to an alternative which supplies too much. The benefit of avoiding moral hazard is more important in the converse case where too much is supplied against Level B.

<sup>48</sup>The PMA allocates the quantities at which the BoE's supply curves meets the bidders' demand curves and it sets prices to implement this allocation. In contrast, in the SSA for example the allocation of funds is fixed for each collateral set so, in a period of stress, demand against Level B might be far higher than the BoE's supply curve at this fixed allocation meaning the allocation of B would be far lower than optimal.

<sup>49</sup>It would be exactly equal if the BoE's supply curve and the bidders' demand curve were linear with equal slopes in absolute value for each good. It is smaller than the product of the two if the average slopes of the supply curves are smaller than the average slopes of the corresponding demand curves in absolute value between the quantities sold in the two designs, the supply curves are weakly convex—which approximates the BoE's supply curves providing the quantity allocated against Level B is not too large—and the demand curves are weakly concave. It is larger than the product of the two under converse conditions.

## 9 Conclusion

We have shown that the flexibility in the allocation of the PMA generates welfare gains relative to the two simple alternatives of the SSA and RPA. The RPA could not have maximised welfare with any fixed reference prices, and hence fixed price difference, due to the instability of the optimal price difference and the lumpiness of bids. And although the SSA could have almost maximised welfare in the period had the BoE been able to accurately forecast the optimal quantities, this result was specific to the stability of the optimal quantities in this particular time period. Moreover, it seems unrealistic to think that the BoE could have forecast these quantities.

The BoE always gained from the PMA relative to the comparators in the period studied. However, the welfare gains of the PMA were not always shared with the bidders. Average bidder surplus was higher in the PMA relative to the comparators only when either the fixed share to Level B used in the SSA was low, or when the fixed reference price of B in the RPA was high. Three features in the period can explain why the bidders did not always gain: the absence of significant stress, only one collateral whose demand is not flat, and little paired bidding. Absent any one of these three features, the bidders, as well as the BoE, might have benefited more from the use of the PMA. Indeed, bidders might benefit more under the current ILTR design relative to the sample period. Since 2014, participants can bid for loans secured by an additional collateral set, which is less liquid than Levels A and B and for which we would not expect demand to be flat. Moreover, the design might have been more beneficial to bidders at the beginning of the COVID-19 pandemic in March 2020, in which demand for loans was acute.

However, our analysis has ignored several significant advantages of the PMA.

One benefit of the PMA relative to the alternatives is that the PMA reduces the scope for strategic behaviour.<sup>50</sup> The PMA combines the markets into a single mechanism, encouraging competition and reducing bidders' incentives relative to the SSA to unilaterally understate their demands. It also may reduce the bidders' ability to coordinate prices relative to the alternative designs, especially when the auctioneer's preferences are not revealed to the bidders, as in the ILTRs. We maintained an assumption of truthful bidding across all designs, eliminating this potential source of welfare loss. If this fails to hold, the actual welfare losses of the SSA and RPA relative to the PMA may in fact be much larger.

---

<sup>50</sup>Grace (in prep., a) tests for strategic behaviour in the ILTR operations and finds that truthful bidding is a reasonable approximation under the PMA in the ILTR auctions.

Furthermore, even in the absence of these effects, the welfare gain of the PMA would likely have been much greater in a period with less stable demand.

Most importantly, neither the SSA nor the RPA would have achieved the BoE’s main objectives in setting up the PMA. Any SSA is simply setting a quantity for funds lent against the less liquid Level B collateral, and any RPA is simply setting a fixed spread between loans secured by collateral assets with different liquidity levels. So neither of these natural alternative mechanisms to the PMA would both allow a substantial increase in allocations of funds against Level B collateral in the event of stressed conditions, but only allow this at higher interest rate spreads, as the BoE wished. Our paper has shown that the PMA enhanced efficiency as measured by the difference between the spreads that financial institutions were willing to pay and the spreads that the BoE was willing to accept on the loans allocated in the auctions. Broader benefits of the auctions, e.g. the ability to stop runs on individual banks (and the externalities of avoiding systemic failure) while avoiding moral hazard, are taken into account, to the extent that the BoE’s supply curve internalises some externalities. Paul Fisher (then Executive Director at the BoE) described the design as “potentially a major step forward in practical policies to support financial stability” (Milnes, 2010). The automatic adjustment to market conditions of both the amount of funds loaned in the PMA, and the interest rate premium charged, improves efficiency not only directly, as analysed in this paper, but also in this broader sense, too.

## References

- Armantier, O., C. A. Holt, and C. R. Plott (2013). “Procurement Auction for Toxic Assets with Asymmetric Information.” *American Economic Journal: Microeconomics* 5(4): 142-62.
- Bank of England (2022). “The Bank of England’s Asset Purchase Facility: APF Operating Procedures; Corporate Bond Purchase Scheme; Gilt Purchases.” <https://www.bankofengland.co.uk/-/media/boe/files/markets/asset-purchase-facility/operating-procedures.pdf>.
- de Roure, C., and N. McLaren (2021). “Liquidity Transformation, Collateral Assets and Counterparties.” *Central Bank Review*, 21(4):119-129.
- Fisher, P. (2011, March). “Recent developments in the sterling monetary framework.” Manchester Economics Seminar, University of Manchester.
- Frost, T., N. Govier, and T. Horn (2015). “Innovations In The Bank’s Provision Of Liquidity Insurance Via Indexed Long-Term Repo (ILTR) Operations.” *Bank Of England Quarterly Bulletin* 55 (2): 181–188.
- Grace, C. (in preparation, a). “Explaining Bidding in the Bank of England’s Liquidity Auctions.”
- Grace, C. (in preparation, b). “Who Gains from Efficient Auctions?”
- Klemperer, P. (2008). “A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions.” Mimeo: Oxford University.
- Klemperer, P. (2010). “The Product-Mix Auction: A New Auction Design for Differentiated Goods.” *Journal of the European Economic Association* 8: 526-36.
- Klemperer, P. (2018). “Product-Mix Auctions.” Working Paper 2018-W07, Nuffield College.
- Lings, S. (2013). “Welfare Effects of Combining Auctions.” M.Phil. Thesis: University of Oxford.
- Milnes, A. (2010). “Creating Confidence in Cash.” Blueprint.
- Nixon, D. (2014). “Sterling Monetary Framework Annual Report 2013–14.” *Bank of England Quarterly Bulletin* 54 (2): 218–225.

Song, Z., and H. Zhu (2018). “Quantitative easing auctions of treasury bonds.” *Journal of Financial Economics* 128(1): 103-124.

Weitzman, M. (1974). “Prices vs. Quantities.” *Review of Economic Studies* 41(4): 477–491.

## Appendices

### A The BoE's algorithms

#### A.1 Detailed description

We simplify by ignoring the exceptional knife-edge cases, that never actually arose, in which the competitive equilibrium prices were not unique, so write the competitive equilibrium prices and quantities as  $p_A^c, p_B^c, q_A^c, q_B^c$ , respectively.<sup>51</sup> Based on the results, and on available information about how the algorithm changed in October 2011 and March 2013:<sup>52</sup>

The BoE's programs always set  $q_A = q_A^c$ .

- (a) From the beginning until September 2011: The BoE set  $p_A$  equal to the highest price for which demand on A equalled  $q_A^c$ . It then computed  $q_B$  based on  $p_A$  and the relative supply curve. It then set  $p_B$  equal to the highest price for B that was consistent with a demand for  $q_B$  (which price could exceed the unique price consistent with a supply of  $q_B$ ).
- (b) From October 2011 until February 2013: The BoE set  $p_A$  and  $q_B$  as in (a) above. It then set  $p_B$  equal to the unique price for B that was consistent with both the demand for, and the supply of,  $q_B$ .
- (c) From March 2013 to January 2014: The BoE chose the competitive equilibrium prices and quantities.

These changes in the PMA price determination rules are summarised in Table 1 in Section 2.

---

<sup>51</sup>If the auction is not fully allocated (in competitive equilibrium),  $p_A^c$  is unique ( $p_A^c = 0$ ), and the monotonic strictly upward-sloping relative supply curve then uniquely determines  $p_B^c$  (at the intersection of the relative supply and the demand curve for good  $B$ ). If, instead, the auction is fully allocated,  $p_A^c$  is uniquely determined except in the knife-edge case in which the numbers of bids taken on A and B exactly “fit” the amount to be allocated without any need for rationing on good A. Even in this case, the difference between the prices,  $p_B - p_A$ , is uniquely determined, since the relative supply curve is upward sloping (see fig. 2 of Klemperer, 2008, for illustration). Moreover,  $q_B^c$  is therefore also uniquely determined. So only if there is also no need for rationing good  $B$ , is there any ambiguity in the prices. (Similarly, in a simple uniform price auction, only if the exact quantity offered is sold without rationing is there a distinction between “HLP” and “LWP”.)

<sup>52</sup>There is less evidence for the period March 2013 – January 2014 than for the other periods, since there were no bids on good A at prices above 0 prior in this period. However, the BoE definitely intended  $p_A = 0$  for non-fully allocated auctions starting in 3/2013, and competitive equilibrium prices and quantities should have been the result. A completely different algorithm was used to find competitive equilibrium prices and quantities from February 2014.

## A.2 Implications for welfare

In all fully allocated auctions, and in non-fully allocated auctions in which  $p_A = 0$ , the results were always as in Klemperer (2008), except in the first period, period (a) above, in which sometimes  $p_B > p_B^c$ .<sup>53</sup> Since the quantities were the same (even in the first period) as in the originally-defined PMA, i.e.  $q_A = q_A^C$ , and  $q_B = q_B^c$ , it follows that the welfare (as conventionally defined) was also the same.

In the two auctions in the first period, (a), and the five auctions in the second period, (b), which were non-fully allocated with  $p_A > 0$ , the results were as in Klemperer (2008) except that  $p_A > p_A^c$  (since  $p_A^c = 0$ ) and sometimes  $p_B > p_B^c$  and sometimes  $q_B < q_B^c$ .<sup>54</sup> However, only in January 2011, December 2011, April 2012, and June 2012,<sup>55</sup> was  $q_B < q_B^c$ , so only in these auctions were the quantities allocated different from Klemperer (2008) and only on good  $B$ . Only in these four auctions, therefore, did social welfare (as conventionally measured) differ from Klemperer (2008).

## A.3 Interpreting the BoE's preferences

The PMA developed for the BoE is described in Klemperer (2008). The proposal was to implement competitive equilibrium allocations and prices, assuming that the BoE's relative supply curve and the counterparties' bids represent their true preferences.<sup>56</sup> (In the event of non-uniqueness, Klemperer (2008) recommended choosing the lowest competitive equilibrium prices,<sup>57</sup> so that small bidders, at least, have no incentive to misrepresent their valuations.<sup>58</sup>) The PMA initially implemented by the BoE was a close approximation to, but differed from, the PMA as originally proposed. It sometimes yielded different results, which create some ambiguity in the appropriate measure of welfare. (Appendix A.4 shows that these differences were infrequent and small, and the BoE subsequently moved in two steps towards an exact

---

<sup>53</sup> $p_A^c$  was automatically identical to the good  $A$  clearing price in the originally defined PMA in fully allocated auctions, except in exceptional "knife-edge" cases that never arose.

<sup>54</sup>It is theoretically possible, though it never happened, that an auction could have been fully allocated in competitive equilibrium but not fully allocated by the BoE in periods (a) or (b). (The reverse could not happen.)

<sup>55</sup>The three auctions of these seven in which  $q_B = q_B^c$  were April 2011, March 2012, and October 2012.

<sup>56</sup>This was justified by the expectation that no bidder could gain much by misrepresenting their preferences.

<sup>57</sup>The lowest competitive equilibrium price vector is uniquely defined—it is the most "south-westerly" point on a 45° line in fig. 1b of Klemperer (2008), if the competitive equilibrium price vector is not unique.

<sup>58</sup>If a bidder expects LWP, they have an incentive to misrepresent their preferences, regardless of their size. They can increase their expected surplus by reducing their bid below their value, which lowers the clearing price, increasing their surplus, in the case in which their bid determines the clearing price. (Large bidders have the additional incentive to misrepresent their preferences by exerting market power, under both the HLP and LWP.)

implementation of Klemperer (2008).)

If the BoE’s relative supply curve represents the minimum price that the BoE is willing to accept to supply good  $B$  relative to good  $A$  (as a function of the quantity of good  $B$  allocated), then the social welfare from any given auction outcome is unambiguous and easy to calculate.

The outcomes of the initial implementation and originally defined PMA differed because choosing competitive equilibrium prices is not always the same as first choosing competitive equilibrium allocations and then choosing (either lowest or highest) prices consistent with total bidder demands equalling those allocations.<sup>59</sup> This is because the auctioneer might prefer different (either higher or lower) allocations at those prices. As discussed in Appendix A.1, above, this difference only arose when cover was low enough that an auction was not fully allocated.

To see this, suppose the only bids are a small bid on good  $A$  at 10bps, and a small bid on good  $B$  at 12bps, and the BoE’s relative supply (of good  $B$ ) for small allocations is 5bps. Then the only competitive equilibrium prices are  $p_A^c = 0$ bps and  $p_B^c = 5$ bps, reflecting the minimum prices that the BoE is willing to accept. With any higher prices on either good, the BoE would not be in equilibrium, since it would wish to sell more than was demanded. A result of this kind arose in the June 2012 auction, in which the BoE’s programme chose  $p_A = 10$ bps and no sale to good  $B$  (as if  $p_B = 15$ bps).

A revealed-preference approach would be that the BoE was in fact always optimising its welfare function. The BoE’s supply function would indicate a preference for a price difference of 5bps in this case, so that if  $p_A = 10$ bps—which would allow it to serve all the demand on good  $A$ —then the correct price for  $p_B$  should be 15bps. The problem is that it is then hard to know exactly what the BoE’s welfare function was. It is unclear why  $p_A = 10$ bps,  $p_B = 15$ bps is superior to  $p_A = 0$ bps,  $p_B = 5$ bps, or other choices with  $p_B = p_A + 5$ . For example,  $p_A = 7$ bps,  $p_B = 12$ bps are the highest prices consistent with total bidder demands equalling those allocations, and might be more “profitable”.<sup>60</sup> And  $p_A = 0$ bps,  $p_B = 5$ bps

---

<sup>59</sup>The BoE’s initial approach was close to choosing competitive equilibrium allocations, and then choosing the highest prices consistent with total bidder demands equalling those allocations (and was always equivalent to this for fully allocated auctions). So for fully allocated auctions the BoE’s initial approach corresponded to lexicographically maximising social welfare as conventionally defined, then, subject to this, maximising BoE surplus. (Klemperer (2008) discouraged this choice of higher than competitive prices, since it gave bidders incentives, albeit probably small ones, to misrepresent their preferences. Hence the BoE’s moves to their subsequent implementations.)

<sup>60</sup>Choosing  $p_A = 7$ bps,  $p_B = 12$ bps would have yielded a greater BoE surplus (assuming the prices that the BoE was willing to accept were 0bps and 5bps) in June 2012.

yield the highest social surplus if we assume the “disutility” (i.e. “cost”) to the BoE of allocating good  $B$  is  $p_A + 5$ . Even if there is a clear reason why  $p_A = 10\text{bps}$ ,  $p_B = 15\text{bps}$  is preferred,<sup>61</sup> it would be unclear what assumption should be made about the BoE’s utility function when making comparisons with the RPA and SSA which won’t generally yield the “correct” price difference.<sup>62</sup>

We have therefore assumed that the BoE all along intended to implement the PMA as originally defined, as is consistent with its behaviour in the latter part of the period studied. That is, we assume that the BoE’s objective throughout was to maximise the sum of bidder surplus and BoE surplus, as defined in Section 3.3.

#### A.4 Relevance of the HLP versus LWP choice to measured welfare

Welfare, as defined in Section 3.3, is only affected by the difference between the BoE’s implementation and Klemperer (2008) when all bids on good  $A$  are allocated, and the lowest bid on good  $A$  is strictly positive, so that the clearing price for good  $A$  differs under the two pricing rules. Under LWP, it is strictly positive and under HLP, it equals zero. In this case the allocation of  $B$  may differ between the two rules (in which case, at least under LWP,  $B$  will not be fully allocated).

Assuming, as in Appendix A.3, that the BoE’s relative supply curve represents the minimum price it is willing to accept to supply good  $B$  as a function of its quantity, the BoE’s actual implementations were slightly suboptimal on a few occasions. Table A.1 shows the actual and optimal amounts allocated, cover ratios and clearing prices in the four auctions in which the allocation in the sample period differs from the competitive equilibrium allocation.

This issue is not economically significant. It affects only four low-cover auctions in the period studied—January 2011, December 2011, April 2012, and June 2012—and the resulting

---

<sup>61</sup>Perhaps the BoE wanted the highest price for good  $A$  that was consistent with demand matching the competitive equilibrium quantity. And perhaps it also had a strong preference for the “correct” spread. The BoE likely did have preferences over the price-spread (independent of the prices that the BoE was willing to accept, as revealed by the supply curves), in order to send an informative signal to the market, and to charge an appropriate premium for loans secured by poor-quality collateral. If, in our example, the small bid on  $B$  was at 5bps rather than 12bps, and if the BoE wanted  $p_A = 10\text{bps}$ , it would have been odd to allocate the bid on  $B$ , which would have required  $p_B < p_A$  but have yielded no welfare gain (based upon the bids and supply curves).

<sup>62</sup>For example, how much worse for the BoE are  $p_A = 10\text{bps}$ ,  $p_B = 12\text{bps}$ , which are the highest prices consistent with total bidder demands equalling the competitive equilibrium quantities? To what extent do the extra allocation of  $B$  (and extra BoE surplus if one assumes the prices the BoE is willing to accept are 0bps and 5bps) mitigate the disutility of the price difference being smaller than the BoE’s relative supply? Note that prior to the October 2011 revision of the algorithm the BoE sometimes chose price differences larger than its relative supply, though it never chose smaller price differences.

Table A.1: Auctions in which the actual PMA is possibly suboptimal in the ILTR auctions, June 2010 – January 2014

Operation date	Amount allocated in actual PMA (% of auction)		Optimal amount allocated (% of auction)		Cover (% of auction)		Clearing spread (bps)	
	Good A	Good B	Good A	Good B	Good A	Good B	Good A	Good B
11 January 2011	54.4	15.65	54.4	17.6	54.4	40	5	26
13 December 2011	16.6	9.85	16.6	11.5	16.6	14.36	3	12
10 April 2012	4	2.7	4	3.7	4	5	4	10
12 June 2012	0.1	0	0.1	2.8	0.1	2.8	10	N/A

measured welfare loss is very small. The difference in net welfare between the SSA and PMA is 0.06bps, 0.06bps and 0.15bps of the actual PMA allocation in these auctions, respectively. Moreover, as discussed above, it is not clear whether this measured welfare loss is real, since there is some ambiguity about the BoE’s preferences in these cases.

A further question then arises as to whether to compare the RPA and SSA with what the BoE actually did, or to what the outcomes would have been if the BoE’s final implementation (i.e. the PMA-HLP) had been used throughout. The latter choice is perhaps the most natural one, since it is the most relevant guide for future policy. In the main analysis, consistent with the other assumptions we have made, we have taken the conservative approach that minimises the average welfare gain of the PMA over the RPA and SSA, and compared the RPA and SSA with the outcomes of the BoE’s actual implementations. That is, we have assumed the BoE made slight mistakes in some early auctions.

Because they were only in a minority of low-cover auctions, they are relatively unimportant and have limited impact on the results. We repeat the surplus difference calculations using HLP for the PMA price determination in all auctions, denoted PMA-HLP. Figure B.7 shows the differences in auction surpluses between the PMA-HLP and SSA, averaged over the June 2010 – January 2014 auctions, weighted by the quantities allocated in the PMA. Figure B.8 shows the analogous surplus differences between the PMA-HLP and RPA.


The net welfare differences are almost identical to the corresponding results for the actual PMA, shown in Figures 3 and 4. The transfers from bidders to the BoE under the PMA are slightly smaller, as the clearing prices are on average lower when using HLP.

Since the PMA price determination rule has been HLP on both goods since March 2013, the ambiguity or potential welfare loss can no longer arise.

## B Additional figures

Figure B.1: Difference in auction surpluses, PMA minus posting pair of fixed prices with Level A price fixed at 0bps, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term





Figure B.2: Difference in auction surpluses, PMA minus posting pair of fixed prices Level A price fixed at 5bps, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term





Figure B.3: Simulated difference in auction surpluses, PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation), transformed bids: Level B quantities demanded multiplied by 8, Level A quantities demanded multiplied by 4, Level B bid prices multiplied by 4

(a) 3-month term



(b) 6-month term





Figure B.4: Simulated difference in auction surpluses, PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation), transformed bids: Level B quantities demanded multiplied by 8, Level A quantities demanded multiplied by 4, Level B bid prices multiplied by 4

(a) 3-month term



(b) 6-month term





Figure B.5: Difference in auction surpluses, BoE-surplus-maximising PMA minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term





Figure B.6: Difference in auction surpluses, BoE-surplus-maximising PMA minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term





Figure B.7: Difference in auction surpluses, PMA-HLP minus SSA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term





Figure B.8: Difference in auction surpluses, PMA-HLP minus RPA, June 2010 – January 2014 (averages, weighted by PMA allocation)

(a) 3-month term



(b) 6-month term

