CHAPTER SEVEN

Bidder Strategies*

I suggest explanations for the apparently puzzling bidding in the British and German 3G telecom auctions in 2000. Relative performance maximization may have been important, but the outcome of the British auction seems to have been efficient.

A. SOME OBSERVATIONS ON THE BRITISH 3G TELECOM AUCTION

7.1 Introduction

As discussed in chapters 5 and 6, and section B of this chapter, the UK auction was one of the most successful of the western European 3G auctions. Indeed in terms of revenue raised per capita it was *the* most successful of all the auctions, and it is therefore appropriate to examine, as Börgers and Dustmann (2002a) (henceforth B-D) do, whether the auction's outcome was also as efficient as is often claimed and whether the bidding in the auction fits well with standard theory.[1]

7.2 Efficiency of the UK Auction

B-D's analysis makes clear that an ascending auction like that in the United Kingdom runs the risk of an at least slightly inefficient outcome arising in some circumstances. However, it also seems clear that the actual outcome of

* This chapter bundles my comments on two papers presented at the December 2001 CES Ifo conference on the telecom auctions, but the chapter can be read independently of those papers. It was originally published as "Some Observations on the British 3G Telecom Auction: Comments on Börgers and Dustmann", in *ifo Studien* 2002, 48, 115–120, and "Some Observations on the German 3G Telecom Auction: Comments on Grimm, Riedel, and Wolfstetter", in *ifo Studien* 2002, 48, 145–156. (This material is also reprinted in *Spectrum Auctions and Competition in Telecommunications*, Illing, G. and Klüh, U. (eds.), MIT Press, 2004.) I am grateful for helpful comments from Tilman Börgers, Elmar Wolfstetter and the representatives of some of the firms involved in the German auction to whom I showed an earlier draft; I am especially grateful to Marco Pagnozzi for our collaboration in the study of the 3G auctions and for his helpful suggestions about these essays.

[1] I learnt a lot from B-D's analysis. In what follows, I discuss just two issues where my interpretation is slightly different. I was the principal auction theorist advising the Radiocommunications Agency which designed and ran the UK auction, but the views expressed in this chapter are mine alone.

the UK auction was efficient, or very close to efficient, in the sense of maximizing the sum of the valuations of the license holders.

Klaus Schmidt's (2002) excellent comment explains that the evidence from the bidding in the auction itself suggests that the UK auction was probably efficient. Evidence subsequent to the auction supports the same claim. It seems clear after the fact—and especially after the other European auctions—that the four incumbents had the highest valuations,[2] so were efficient winners. There is no evidence that any losing entrant had a value for a license that exceeded TIW-Hutchison's. Finally, the evidence subsequent to the auction, as well as from within it (including the interpretation of the bidding offered below), suggests that Vodafone had a higher incremental value for a large license than did any other incumbent, and therefore that the allocation of licenses among winners was also correct.

In short, all the available evidence suggests that the UK auction's outcome was efficient in the sense claimed.[3]

7.3 BT's Bidding Behavior

B-D also suggest that some of the behavior they document is very hard to rationalize, but I conjecture that doing sufficient research into the environment in which the auction took place will yield good explanations; I illustrate this by examining the main "puzzle" of BT's bidding.[4]

BT's bidding was such that the prices bid for the large (2×15 MHz) "B" and small (2×10 MHz) "C", "D" and "E" licenses differed by roughly a constant in the early stages of the auction (phase 1 of the auction in B-D's terminology), and then switched to differing by roughly a fixed proportion (50 percent of the price level of the small licenses) in the later stages of the auction (phases 2 and 3 in B-D's terminology).[5] This pattern seems unusual, but reviewing analysts' reports provides a clue: some analysts assumed the value of the large license must be 1.5 times the value of a small license (reflecting an assumption that 1.5 times the amount of spectrum would

[2] See van Damme (2002) and Fortis (2000) for evidence and discussion of these value differences. (Indirect evidence is also provided by the fact that only one out of the 30 incumbent bidders in the eight western European ascending auctions failed to win a license—and even this single failure was attributed to collusion or organizational strife within the bidder, rather than to the incumbent having a low value, see section 5.3.3.)

[3] Cable, Henley, and Holland (2002) use stockmarket data to argue that 'there is no evidence that the outcome of the auction was anything but efficient'.

[4] However, B-D are to be congratulated on having already explained so much; they also looked at evidence from outside the auction to explain behavior within it.

[5] That BT's behavior in the later stages of the auction can be described in this way was observed independently, by B-D and myself, after the conference in Munich to which their paper was contributed. The details are reported in B-D's companion paper, Börgers and Dustmann (2002b).

allow 1.5 times the service to be offered[6]), while several other analysts insisted that the large license was worth a fixed sum more than a small one (reflecting the additional costs—base stations, etc.—required to run the same service with a smaller license), and it was clearly well understood in the industry that different bidders might make different choices between these two different valuation models.

Of course, if one or more bidder valued the large license at 1.5 times the value of the small license, this cannot on its own explain the price difference being a fixed proportion of the value of the small license. For example, if BT's private valuations for small and large licenses were £4 billion and £6 billion, respectively, while Vodafone's were £6 billion and £9 billion, respectively, and other bidders were closer to indifferent between small and large licenses, then with "straightforward bidding" (in B-D's terminology) the absolute value of the price difference would quickly move to equal £2 billion (since whenever the price difference was less than £2 billion, both BT and Vodafone would regard the large license as the best deal, and so would bid on it).[7]

However, it seems plausible that BT intrinsically valued a large license more than a smaller license by a fixed value that was considerably below 50 percent of the final price of a small license. BT may also have become very confident that Vodafone valued a large license at 50 percent more than a small license. (Apart from any information from outside the auction, Vodafone never placed a bid on any license other than the large license in the auction.) Furthermore, BT may have wished to make Vodafone pay as much as possible for its license[8] for at least two reasons. First, this would reduce Vodafone's budget and so make Vodafone a weaker competitor in subsequent auctions (the British auction was the first of nine western European 3G auctions, and was also followed by others elsewhere in the world). Second, making Vodafone pay more would make "the market" think Vodafone had not done better than BT in the auction. There is anecdotal evidence that BT was very concerned both about the stock market's perceptions of its performance, and about the wider market's view of its position relative to Vodafone. Allowing Vodafone to win the larger license at a lower per MHz price than BT was paying might suggest BT's managers had got a bad deal. Or it might suggest that BT was not able to make effective use of a larger license in the way that Vodafone could, and hence that BT thought it was in a weak market position,

[6] The technology might actually allow offering slightly more than 1.5 times the service, hence the value ratio might be slightly more than 1.5.

[7] Even if, as I will argue, some of the early bidding was non-serious, the price difference would move to the fixed amount, £2 billion, as soon as the bidding became serious.

[8] After the auction BT claimed it had deliberately pushed up the price that Vodafone had paid, and this was reported in the press (see Cane and Owen 2000). (At the time, this claim was dismissed by auction theorists as implausible, since it was hard to reconcile with the evidence without realizing that BT and Vodafone might both have had different valuation models and also have had a reasonably clear idea of the other's valuation model.)

while Vodafone was clearly "number one".[9] So bidding up the large license's price to 50 percent more than the current small license price may have seemed a reasonable risk to take, even given the small chance of ending up winning the large license at hundreds of millions of pounds more than BT valued it.[10]

Of course, even a small risk of winning the large license might seem to have a significant expected cost. But it was also possible that if BT did end up winning the large license, it might have been able to resell part of it at little or no loss, given that the auction prices would then have established a clear price per MHz. (The possibilities for resale were unclear, but Hutchison did in effect resell a fraction of the license it won, very shortly after the auction, to KPN and Docomo at almost exactly the price per MHz that BT and Vodafone paid in the auction.[11]) In any event, observers might not think BT's managers had made a bad decision, even if BT did end up winning (and keeping) the large license for 50 percent more than the price of a small license.[12]

This theory leaves an important question unanswered. Why did BT not push up the price of the large license in the early stages of the auction? One reason is that much of the bidding in the early stages of the contest, when it was clear that there was no realistic chance of the auction ending very quickly (B-D's phase 1) does not seem to have been very serious.[13] In fact, some bids were probably slightly frivolous, or designed to attract media attention. For example, One-2-One raised its bid by slightly more than the minimum required in round 76 to bid £1,212,100,000![14] And BT did start pushing up the price difference between the large and small licenses in round 99 when there

[9] Section 7.8 and Abbink et al. (2001b, section 4) discuss the importance of bidders' concerns about relative performance in two other auctions in which BT's and Vodafone's subsidiaries both competed, the German 3G auction and the previous year's German DCS-1800 spectrum auction. (Strictly, Vodafone was not involved in the earlier auction, but Mannesman, which was a subsidiary of Vodafone by the time of the UK auction, did compete in the earlier auction.)

[10] If BT was correct in its assessment that Vodafone's valuation of a large license was (at least) 50 percent more than that for a small license, the (only) risk that BT faced was that Vodafone would quit the auction altogether. But this outcome was completely implausible, since it would imply that Vodafone's valuation for a small license was below that of Orange and One-2-One (which were both weaker incumbents) and at least one new entrant. The real risk would have been that BT had misjudged Vodafone's valuation difference between the licenses, and BT perhaps knew this risk was small.

[11] The UK Government now seems likely to make resale relatively easy, but this was unclear at the time of the auction, and actual resale of part of a license may in any case be unattractive since bringing a new competitor into the industry makes the remaining spectrum less valuable. Bringing new partners into a joint venture as Hutchison did therefore seems the most relevant form of resale.

[12] Of course, the arguments of this paragraph are in effect postulating that there may have been important common value elements to valuations. Note that with common value elements, it is plausible that the large license might be worth a fixed amount (say £500 million to £1 billion) more than a small license at low prices, but a constant fraction (say 150 percent) of the small license at large prices.

[13] Four bidders have informally confirmed this.

[14] Additional 1's and 2's were ruled out, because all bids were required to be multiples of £100,000.

were still nine bidders left (so four more dropouts were still required to end the auction), and did not then stop pushing up the price difference until round 112 when the large license was more than 50 percent (and more than £1.5 billion) more expensive than the small licenses.

A more serious reason why BT did not push up the price difference earlier is that BT may not have wanted to influence other bidders too early to think that license values were very high (since these other bidders might need time to adjust their views, and get extra money approved by their Boards, etc.). For example, if BT's valuation for a small license was £5 billion, it might have been confident that Vodafone's value exceeded £4 billion for a small license, and therefore that Vodafone would pay at least £2 billion more for a large license. But pushing the price difference up to £2 billion immediately would have sent a very clear signal about what the ultimate prices might be at a time at which the auction prices for the smaller licenses were still very low, and this could only have been damaging to BT's interests.

A final possible reason why BT did not push up the price difference early on is that BT may not have become confident that Vodafone's valuation of the large license was 1.5 times its valuation of the small license until later in the auction.

Most likely BT thought that the early bidding was probably not very important but that its best strategy was to roughly mimic what straightforward bidding would have been if it had had low valuations and a correspondingly low difference in valuations. Certainly this is consistent with the evidence.[15]

So it seems possible to give a reasonable explanation for BT's bidding. Of course, this may not be the only possible explanation.[16] However, the moral is that understanding bidding in auctions often requires knowing a lot of real-world detail about the players and the context in which they are operating. Facts from outside the bidding itself—in this case knowing the differing valuation models that different analysts used—may be the key to explaining behavior. In understanding auctions, as well as in designing them, "the devil is in the details".[17]

B. SOME OBSERVATIONS ON THE GERMAN 3G TELECOM AUCTION

7.4 Introduction

The German 3G spectrum auction was undoubtedly a success from the government's viewpoint. Indeed, it was probably one of only three successes among the nine western European 3G auctions. The measure of success most

[15] Although Vodafone only bid on the large license, it is very plausible that Vodafone was following a similar strategy, but mimicking a bidder with slightly lower valuations.

[16] For example, there may have been much stronger common value components to valuations than usually assumed.

[17] See chapters 3, 4, and 5 for more discussion of the importance of understanding the wider context, and of apparently small details, in auction design.

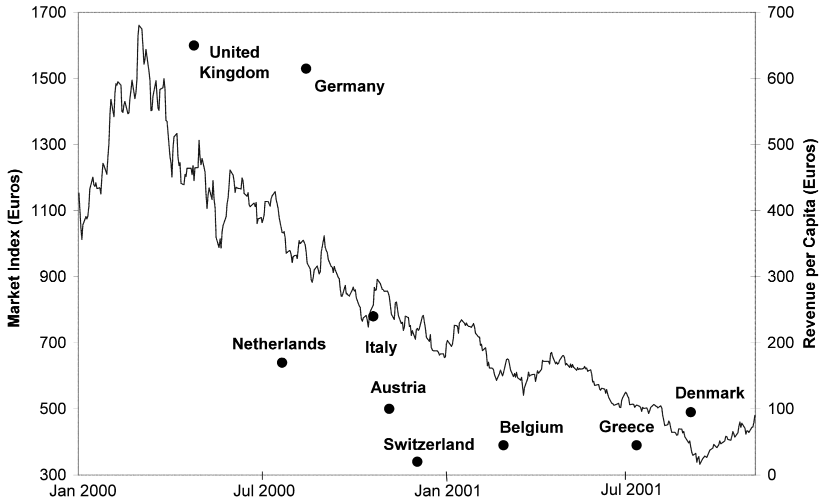

Figure 7.1 European 2000–2001 3G Mobile Spectrum Auctions.
Note: Per capita revenues, by country ● are given by the right-hand scale; auctions are shown on the dates at which they finished. Dow Jones European Telecom Stock Price Index is given by the left-hand scale.

commonly used is total revenue raised per capita, with some adjustments for the level of the telecoms stock index as a reflection of sentiment towards 3G's prospects.[18] (We assume governments have no ability to time the market, and therefore deserve neither credit nor blame for selling when market sentiment is unusually positive or negative.) Based on this, the figure suggests the UK, German, and Danish auctions were successes, while the Dutch, Austrian, and Swiss auctions were the biggest failures (figure 7.1).

Figure 7.1 flatters larger countries (especially Germany, conversely it underrates tiny Denmark), flatters centrally located countries (Germany, again, and also Austria and Switzerland), flatters countries with lightly regulated telecom industries (Germany, again, among others)—since larger, centrally located, lightly regulated markets are worth more—but it also ignores the fact that Germany and Austria sold more licenses than other countries, reducing the total profitability of those markets. However, the more systematic discussion of the relative performance of the different auctions in chapter 5 comes to very similar conclusions.

Since the German auction was a success and was of a novel and complex design, it clearly deserves study, and Grimm, Riedel, and Wolfstetter's (2002) paper (henceforth GRW) would be welcome for that

[18] Although efficiency was generally the primary objective, there is no evidence that efficiency differed much across the different countries' auctions; hence the focus on revenues.

reason alone; their paper gives very valuable detail about the auction and will be a key reference for anyone studying it. But more than that, the paper is extremely interesting and makes acute observations about both the German and other 3G auctions.

I have learned a lot from the paper, and agree with much of it. However, these comments will naturally focus on the disagreements. Section 7.5 summarizes why I think GRW's explanation of the bidding in the German auction is not fully satisfactory, and section 7.6 develops this point more fully (the latter section can be omitted by readers who do not want too much detail). Section 7.7 suggests other explanations for the bidding, and section 7.8 develops a relative performance maximizing theory for it. Sections 7.9 and 7.10 briefly comment on some other 3G auctions, and stress (as GRW also do) the importance of attracting entry into an auction.

7.5 The German Auction

In particular, I disagree with GRW's central claim that their model, as it stands, rationalizes the behavior of the two strongest bidders, T-Mobil and Mannesman, or "T" and "M", in GRW's terminology.[19] These bidders initially pushed up the price in the hope of driving out the sixth-strongest bidder, "bidder 6" in GRW's terminology, but then gave up pushing the price up so that the auction did actually end with six winners but at a much higher price than was necessary to end the auction with this number of winners. This seems bizarre. To put the point simply, consider T's and M's decision about whether to end the auction with six winners at some given price, or whether to push the price up further. Raising the per-block price by 1 euro costs T and M 2 euros each, since they would each win two blocks in a six-winner outcome. Their gain is the probability that bidder 6 quits, times their benefit from bidder 6 quitting. If it is worthwhile for T and M to push the price up in one round, but to stop pushing the price up in the next round, then the perceived probability of bidder 6 quitting in the next round must be both low, and also lower than it was in the last round. However, most observers thought the probability of bidder 6 quitting in the next round, conditional on not having previously quit, was high and increasing around the time T and M ended the auction (when per capita price levels were approaching those achieved in the United Kingdom) and was much lower earlier (a six-player conclusion for the auction became possible at 55 percent of the UK price levels). So any rationalization of T and M's behavior must explain this apparently irrational behavior of theirs. But GRW's model side steps this basic issue, as I now explain.

[19] T-Mobil and Mannesman were subsidiaries of Deutsche Telekom and Vodafone, respectively.

7.6 GRW's Analysis of the German Auction

[**This sub-section can be omitted by readers who do not want too much detail.**]

To understand GRW's argument—and why I believe it is incomplete in this context—consider the preferences of either one of the two strong bidders, M and T. At any point in time, it would like to end the auction only if it prefers this to waiting until the price has risen a small further amount, $\Delta price$, before the auction ends.

The gain from waiting is the probability, ρ_6, that bidder 6 will quit in the price interval, $\Delta price$, times the value of driving bidder 6 out. This value is itself the benefit, β, of winning a third block, including the benefits from excluding bidder 6 from the industry (leading to a more concentrated, and hence more profitable, market), less the current price of buying an additional block. That is, the total gain from waiting equals $\rho_6[\beta - price]$.

The cost of waiting is the extra price, $\Delta price$, paid on the two units that the bidder will win anyway, that is, $2(\Delta price)$. That is, the bidder would prefer to plan to end the auction if a further price rise of $\Delta price$ fails to drive out bidder 6, rather than end it now if

$$\rho_6[\beta - price] > 2[\Delta price]. \tag{1}$$

The bidder would prefer to end the auction now if

$$\rho_6[\beta - price] < 2[\Delta price]. \tag{2}$$

In GRW's model, bidder 6's value can only take one of two possible values, v_6' (strong type) and v_6 (weak type), and the auction cannot be ended before price p ($< v_6$), so the only conceivably sensible strategies for the strong bidders are:

(a) to end the auction as soon as possible at p, or
(b) to push the price up a further $\Delta price = v_6 - p$ to v_6 to drive out the weaker type of bidder 6, but then to end the auction, or
(c) to push the price up further still, by an additional $\Delta price = v_6' - v_6$ more to v_6', to drive out both types of bidder 6.

The observed behavior in the actual auction corresponded to case (b) of GRW's model.

The condition for (b) to be preferred to (a) is the appropriate version of (1) or equivalently is GRW's equation (9).[20] The condition for (b) to be preferred to (c) is the appropriate version of (2), or equivalently GRW's equation (6).[21] So these are the key conditions in GRW's Theorem.[22] If (6) holds there is an equilibrium in which outcome (b) arises. (If (6) fails, both strong bidders prefer (c), and either

[20] In GRW's notation, $\beta = v_{13} + b$ for the case of successfully driving out bidder 6 at price v_6.

[21] In GRW's notation, $\beta = v_{13}' + b$ for the case of driving out bidder 6 at price v_6', but $\rho_6 = 1$ in GRW's model at stage (c).

[22] GRW rename (6) as "$\Delta_{pr}^2 \leq 0$" in their statement of Theorem 1.

can unilaterally implement it.) If (9) holds, the equilibrium is unique. If (9) fails, outcome (b) can still be an equilibrium of the model if (6) holds, since neither strong bidder can unilaterally end the auction, but the equilibrium is neither unique, nor plausible.[23] So for the observed play to correspond to a plausible equilibrium of GRW's model, both (9) and (6) are required.[24]

Noting that (9) and (6) are just my equations (1) and (2) suggests why GRW's theory seems unlikely to describe reality. Of course, (9) and (6) correspond to (1) and (2) evaluated at different values of ρ_6, β, *price* and Δ*price*, reflecting the different stages of the game at which (9) and (6) are computed. So the observed play can correspond to an equilibrium of GRW's model. But this requires that $[\rho_6/(\Delta price)]$ not be too much lower when the strong bidders could first have ended the auction (when (1) must be satisfied) than at the actual end of the auction (when (2) must be satisfied).[25]

Furthermore, the tension between conditions (9) and (6) is more severe when the game is generalized to many small stages since the values of $[\rho_6/(\Delta price)]$, β, and *price* to be substituted into (1) and (2) cannot then vary much between stages, and related conditions must then hold at all the stages—an issue that GRW do not address.[26] In particular, (1) must hold just

[23] This equilibrium is not plausible if (9) fails because in this equilibrium both strong players prefer outcome (a) to outcome (b), but both follow the strategy corresponding to (b) because each expects the other to do this. This logic can only hold in the two-stage model: with more stages, each strong player would know that if it followed the strategy corresponding to (a), then the other strong player would follow just one round of the auction later (if (9) fails)—that is, the players could trivially coordinate on strategy (a) in the actual multi-round auction, which eliminates this equilibrium. This equilibrium is, of course, also Pareto dominated by the more natural equilibrium for the players in GRW's two-stage model, and GRW also eliminate this equilibrium in their limited extension to multiple rounds (see note 26). However, an equilibrium of this kind becomes more plausible if M and T are each uncertain that its rival shares its assessment of the parameters, or are uncertain about the rival's objectives (see section 7.8).

[24] GRW note that Theorem 1 also requires other conditions.

[25] It does not seem likely that $(\beta - price)$ ever became very small because β includes both the value of a third block to a strong bidder and the value, b in GRW's terminology, of excluding bidder 6 from the industry, leading to a more concentrated and hence more profitable market. Therefore, β must be greater than $p_2 + b$, where p_2 is the expected maximum of the value of a fourth block to a strong bidder and a third block to a less strong bidder. And, as GRW point out, p_2 must itself be quite high for the GRW equilibrium to make sense—the logic of GRW's equilibrium requires p_2 to be at least equal to the final German auction price. (The very limited anecdotal evidence suggests that p_2 might have been, very roughly, in the region of the final German auction price.)

[26] GRW do briefly consider extending their model to many rounds of bidding, but when they do this they maintain the extreme assumption that bidder 6's valuation can take only two possible values, v_6 and v'_6. Thus, in their extension there is no possibility of bidder 6 quitting before v_6, or between v_6 and v'_6, so the additional rounds of bidding are mostly irrelevant and (1) and (2) are relevant only at the same points at which they matter in the two-stage game, i.e., GRW's conditions (9) and (6) suffice as before. In a proper many-round extension of GRW's game in which it is also recognized that bidder 6's valuation is not restricted to just two possible values, conditions related to (1) or (2) must hold at each round of the game. (One difference that arises even in GRW's simplified many-round version is that (9) and (6) are both required for GRW's result to be an equilibrium.)

before the auction ends, and (2) must hold at the price at which the auction ends. If the price is changing only slowly between rounds (as was the case in the actual auction), it is required that $[\rho_6/(\Delta price)]$ is falling (or at least not increasing much) at the end of the auction.

Summarizing the two previous paragraphs, GRW's equilibrium requires that, at the end of the auction, the probability of bidder 6 quitting conditional on not yet having done so is *both* not much increasing, *and* not much larger than at the lower prices at which the strong bidders could earlier, if they had both wished, have ended the auction.

These two conditions seem implausible. A six-player conclusion to the auction became possible when Debitel quit at prices that were just 55 percent of the final UK prices (per capita).[27] The German auction actually finished at 94 percent of the final UK prices. The weakest of the six remaining bidders was generally thought to be either Mobilcom or "Group 3G", the joint venture between Telefonica and Sonera, so ρ_6 represents the probability that one of these would quit in the next round, conditional on their not yet having quit. But Telefonica and Mobilcom had quit the UK auction when the price levels reached 94 percent and 100 percent of the final UK price level, respectively.[28] Mobilcom (at least) had made public statements that suggested that it was likely it would bid as far as it had in the UK auction,[29] and outside observers also thought that these bidders would probably go a lot further than 55 percent of the UK auction price, but might quit at around the final UK price levels. Certainly, most plausible distributions of valuations implied that at the end of the auction the probability of bidder 6 quitting was both much higher than earlier, and increasing, and either of these implications is sufficient to rule out GRW's equilibrium.[30]

In brief, GRW's equilibrium requires, roughly, that the strong bidders thought it relatively likely that Mobilcom or Group 3G would quit while prices were well below UK levels but then, having seen that Mobilcom and Group 3G did not quit at such low prices, the strong bidders thought it both relatively

[27] Ending the auction at these prices would have required the cooperation of all six bidders, but this could probably have been obtained. Even if this seemed hard, M and T could together have ended the auction once the other four bidders had stopped bidding for three blocks; none of the other bidders had high bids for three blocks beyond round 136 when the prices were 70 percent of the final UK price, and all could be proved to have lost eligibility for three blocks shortly thereafter.

[28] Mobilcom was owned in large part by France Telecom which was also part owner of NTL Mobile, the last bidder to quit the UK auction.

[29] It may be said that such statements were cheap talk. But following through on them may be necessary to maintain management credibility; they probably reflected an availability of finance, and—what matters—they seemed credible to observers at the time.

[30] Although GRW argue (in the last paragraph of their section V) that it was reasonable to expect that bidder 6 might quit at some point before the final German auction price, they fail to consider the crucial questions about the relative likelihoods of bidder 6 quitting at a very low price, or at close to the final UK prices.

unlikely, and increasingly unlikely, that they would quit while prices were close to UK levels. This seems unreasonable.

7.7 What Actually Happened in the German Auction?

Whilst no one can be certain, it seems that other factors are required to explain the behavior of T-Mobil and Mannesman in the German auction. Some of these factors are discussed in sections 4.5 and 5.4.1.[31]

They include the complexity of the rules and the opacity of the information available to bidders about others' bids, which made it hard for bidders to figure out optimal strategies (T may simply have made a mistake in failing to heed M's signal suggesting that they both reduce demand early on) or to understand their rivals' thinking. Section 4.5 stresses the apparent lack of trust and understanding between the two strong bidders, and discusses why this mistrust might have arisen.

Furthermore, the strong bidders may not simply have been maximizing expected profits. M and T may have focused more on their performances *relative* to each other, as might be rational behavior for managers who had private career concerns, or were concerned that their firm seemed well managed and deserving of further investment, etc. Relative performance concerns may explain the auction's outcome, especially in conjunction with the mistrust between the bidders, as we explain in more detail in the next section.

Other contributory factors to T's behavior that have been suggested include that T felt pressured by the stock market's response to the rising auction prices (and that T had not fully anticipated this), and even that T's objectives were affected by the fact that it was majority owned by the German government.

7.8 A Relative Performance Maximizing Theory of the German Auction

GRW explain that if, for example, M reduced demand to two blocks while T did not, and T then won three blocks by driving out bidder 6, there would then have been a second auction for the remaining block which would most probably have been won by M at an expected price p_2 (in GRW's terminology) so both M and T would have ended up with three blocks but having paid different prices for them.

[31] Ewerhart and Moldovanu (2001) make interesting points about the German design but in a model in which there is only a single strong bidder, so they cannot address why initially both strong bidders pushed up the price, and then both stopped doing so. Also they do not model the second auction that would have taken place if just one strong bidder had pushed up the price and subsequently driven a weak bidder out, and this possible second auction may have played an important role in the behavior in the main auction, as we discuss in section 7.8.

Recall also from our discussion above that when prices are still low (e.g., around 55 percent of the final UK auction price) the probability of bidder 6 quitting is low, so it probably maximizes both firms' expected profits to reduce demand to two blocks and end the auction at low prices. However, if one firm, say M, reduces demand while T fails to do so and continues to push the price up, there is some—perhaps small—probability that bidder 6 will be driven out at a price $\tilde{p} < p_2$, in which case T and M will both end up with three blocks (assuming that M wins the block in the second auction), but T will on average pay less for its blocks than M (since T pays $3\tilde{p}$, but M expects to pay $2\tilde{p} + p_2$). Even in this case T and M may both be worse off in absolute terms than if T and M had both reduced demand to win two blocks at low prices. Because the chance of driving out bidder 6 at a low price is not that high, the more probable result would simply be that T would eventually reduce its own demand to two blocks later on, in which case both T and M would be much worse off than if they had both reduced demand earlier. But note that T always improves its performance relative to M by failing to reduce demand at prices below p_2.

Furthermore, even if each firm is actually an ordinary profit maximizer, but each firm expects that the other is likely to maximize relative performance, then neither firm will reduce demand first (since being the only firm to reduce demand when prices are low risks paying $2\tilde{p} + p_2$ rather than $3\tilde{p}$).[32]

Similarly, when prices are higher (e.g., close to the final UK levels), it may maximize both M's and T's expected profits to push up the price to drive out bidder 6. But if one of the firms, say T, reduces demand to two blocks and lets M push up the price on its own to drive out bidder 6 at a price $p^* > p_2$, then again T and M will both end up with three blocks (assuming that T wins the block in the second auction) but T will pay less on average for its blocks than M pays (since T expects to pay $2p^* + p_2$, but M pays $3p^*$. So T would improve both its relative and its absolute performance if it could reduce demand alone, and M would then improve its relative performance by reducing demand along with T, even though M might increase its (and T's) absolute profits by continuing to raise the price to drive out bidder 6.

The story told thus far is extreme. True, there is anecdotal evidence that firms' managers cared about relative performance, and concerns about relative performance also seem to have played at least some role in other European 3G auctions (see, e.g., section 7.3).[33] But M and T were surely not concerned only with relative performance. So one might have expected M and T to attempt to coordinate their behavior to reduce their demands at low prices to maximize both of their absolute profits. Indeed it seems that M did initially try to signal to T that they should do

[32] A similar argument is that if all firms are known to be ordinary profit maximizers, but firms are unsure that their rival has the same estimates of parameters such as ρ_6, then firms may be unwilling to reduce demand first.

[33] Abbink et al. (2001b, section 4), provides some evidence that relative performance issues were important to these bidders in the German DCS-1800 auction.

just this (see GRW and sections 4.5 and 5.4.1). But T could not know whether M was sincere, and the firms apparently mistrusted each other's intentions (see section 4.5) and, as we have seen, there are very strong relative performance arguments (it suffices that each feared that the other might maximize relative performance) why neither was prepared to be the first to reduce demand while prices were still low.

T then reduced demand later when prices were higher, perhaps for relative performance reasons,[34] and/or because this could also improve its absolute performance if M failed (or was unable) to follow its demand reduction.[35,36] Once T had reduced demand, there are several possible reasons why M followed straight away. First, M had a strong relative performance incentive to follow immediately, as explained above. Second, M may have wanted to develop a reputation for cooperative behavior in which M and T parallel each other's behavior—a kind of 'relative performance' effect but strictly driven by M's long-run absolute performance goals (see section 4.5). Third, M might have been concerned only with its (short-term) absolute performance, but it might all along have taken the view that this would be maximized by M and T both reducing demand, and it might have stuck to this view (i.e., M may have been extremely pessimistic about driving out bidder 6 at low prices and, even though driving out bidder 6 seemed more likely at high prices, remained fairly pessimistic—see section 5.4.1); this is consistent with M's early behavior (signaling T to reduce demand but not unilaterally reducing demand) if it feared that T might place a large weight on relative performance.

[34] It might seem that a firm could protect its relative performance by following a strategy of quitting only if its rival quits when prices are above p_2. However, it takes time to be sure the rival has quit (because the auctioneer gave the bidders only limited information about their rivals' bidding), and it also takes time to respond. Furthermore, some of the weaker players may have been staying in the auction in the hope of being a winner in a five-firm industry, which would have been the outcome if M and T had successfully driven one of them out—in particular, each of Mobilcom and Group 3G might have hoped that the other (or possibly E-plus or Viag) was the 'bidder 6' who might have been driven out. In this case, when one of M and T quits bidding for a third block, these weaker players may expect the other of M and T to try to follow suit and may therefore try to quit first rather than find themselves stuck as winners in a much less profitable six-firm industry. So if either M or T failed to quit first when prices became high, it might have risked being stranded buying a third block at a higher price than its opponent, and achieving a very poor relative-performance.

[35] If T thought M was not too concerned with relative performance, T could improve both its relative and absolute performance by reducing demand once the price exceeded p_2, and free-riding on M continuing to push price up to drive out bidder 6. Even if M was concerned with relative performance, there was the possibility that M would have been unable to follow T (see note 34).

[36] Of course, there may be other reasons, such as T being influenced by stock market pressure or its government ownership (see section 7.7). It might be argued that another possibility was that β was not in fact that high. However, this seems less likely since β must have substantially exceeded p_2 (see note 25) and if p_2 were low then both firms would have been willing to reduce demand early on for relative, as well as absolute, performance reasons.

Of course, there may be other reasons for the observed behavior in the auction. For example, fear that one's rival has very different perceptions from one's own about the chance of driving out bidder 6 can have similar effects to fear that one's rival is a relative performance maximizer,[37] and section 4.5 emphasizes the mistrust and misunderstanding between the bidders. But the point is that the apparently puzzling behavior can be explained by postulating only a limited concern with relative performance. To explain why M and T failed to reduce demand early on, it suffices that each firm thought its rival put some weight on relative performance; it is not necessary that either firm actually did so, and even the conjectured weights on relative performance need not have been large if firms were also uncertain about their rivals' perceptions about bidder 6's behavior, etc. Not much more concern with relative performance is needed to explain the firms' later behavior in the auction.

7.9 The Austrian and Swiss Auctions

Turning to other 3G auctions, I disagree with GRW's assertion that the Austrian auction design was superior to the Swiss, except to the extent that the Austrian reserve price was somewhat more realistically chosen than the Swiss reserve. Neither auction attracted more bidders than there were winners, and neither involved any significant bidding. (Although there was a semblance of serious bidding in the Austrian auction, the bidders there were put under considerable pressure from the authorities to continue the bidding, and it was widely believed that the bidding only lasted the few rounds it did in order to create some public perception of genuine competition and reduce the risk of the government changing the rules.) Neither auction achieved more than 11 percent more than the reserve price that had been set. The only important difference is that the Swiss reserve price had been set ludicrously low at 20 euros per capita, while the Austrian reserve price, although still far lower than it should have been, was 90 euros per capita.[38] But revenues in excess of 300 euros per capita should probably have been attainable in both auctions (see section 5.5). So both of these auctions were failures, and both were intensely embarrassing to their respective governments. Indeed there was no successful European 3G auction after the UK and German auctions until the Danes switched to a sealed-bid design. I have discussed all these auctions in more detail in chapter 5.

[37] In particular, fear can make a bidder unwilling to reduce demand first when prices are low, because of the perceived risk that the rival will not follow (see note 32).

[38] Of course, Switzerland sold four licenses while Austria sold six, but the Swiss could obviously have used the same design to sell six licenses if they had preferred that outcome.

7.10 The Importance of Entry, and the UK Auction

Where I do agree very strongly with GRW is on the importance of attracting entry into an auction.[39] As GRW say, "competition is not a free good" and auctions must be designed with this in mind. However, this does not imply that there is any single best design. Often a sealed-bid design is best for attracting entry, as is suggested by the Danish example in the previous paragraph. But this need not be the case. The UK design was appropriate in its context, because the UK auction was the first 3G auction and was therefore unlikely to suffer from entry problems. (See sections 4.5, 5.3.1, and 5.7 for more discussion of why being first was so important.) Indeed the UK auction attracted 13 bidders compared with the seven that entered the German auction. It seems improbable that the German design would have usefully increased competition in the British auction, and the British design had other advantages over the German design,[40] see chapter 5 and section 6.4.3 (though my view may be colored by my having been the principal auction theorist for the UK auction[41]). In another context, when Peter Cramton, Eric Maskin, and I advised on the UK's 2002 auction for greenhouse gas emission reductions, we chose a uniform price ascending design as being most likely to attract "small" bidders who did not have the resources to work out how to bid correctly in a discriminatory price auction (see Klemperer et al., forthcoming).[42] Nor, of course, is entry always the key issue. As I discussed in chapters 3–6, good auction design is not "one size fits all", but must always be tailored to its context.

[39] I emphasized this in Bulow and Klemperer (1996), Klemperer (1998), and chapters 3–6.

[40] One advantage is identified in note 34: in the German design a bidder might rationally follow a strategy that could mean that it felt sorry to have won as soon as the auction finished.

[41] I was the principal auction theorist advising the Radiocommunications Agency which designed and ran the UK auction. Ken Binmore led the team and supervised experiments testing the proposed designs. Other academic advisors included Tilman Börgers, Jeremy Bulow, Philippe Jehiel and Joe Swierzbinski.

[42] Larry Ausubel and Jeremy Bulow were also involved in the implementation of this auction. This was strictly a descending auction, since the auctioneer was buying reductions in emissions rather than selling permits to emit, but the auction corresponded to an ascending auction to sell emissions.