A1. (i) See section 1.6, section 4.2.1 page 128, and Appendix 1.C.

(ii) [Appendix 1.D pp. 55-57 contains a common-values example closely related to the private-values example described here.]

(a) Bidders bid up to their value, so the price is the expected second highest actual value, i.e. $\theta + \left(\frac{n-1}{n+1}\right)$ (see Appendix 1.E). Profits are $\left(\frac{n^2-n}{n^2+n}\right)$.

(b) Each bidder bids $v_i - x$ for some x (since the diffuse prior means she has no knowledge of how high or low her value is relative to others).

A bidder who deviates to bid $v_i - x + \epsilon$ would earn $x - \epsilon$ if she won, and wins with probability $(t_i + \epsilon)^{n-1}$ and, so she would expect to earn

$\pi(\epsilon) = \int_{t=0}^{1} (x-\epsilon)(t+\epsilon)^{n-1} dt = (x-\epsilon)\left[\frac{(t+\epsilon)^n}{n}\right]_{t=0}^{1} = \frac{(x-\epsilon)}{n}\left[(1+\epsilon)^n - \epsilon^n\right] \Longrightarrow$
$\frac{\partial \pi(\epsilon)}{\partial \epsilon} = -\frac{[(1+\epsilon)^n - \epsilon^n]}{n} + \left(\frac{x-\epsilon}{n}\right)\left[n(1+\epsilon)^{n-1} - n\epsilon^{n-1}\right].$
In equilibrium $\frac{\partial \pi(\epsilon)}{\partial \epsilon} = 0$ at $\epsilon = 0$, i.e. $x = \frac{1}{n}$.
The highest bidder (with expected value $\theta + \frac{n}{n+1}$) wins so profits are $\frac{n}{n+1} - \frac{1}{n} = \left(\frac{n^2-n-1}{n^2+n}\right)$, confirming part (i)'s result that sealed-bid profits are below ascending profits when bidders' private signals (here v_i) are affiliated.

(iii)(a) Ascending auction behaviour is unaffected.

(b) If θ is common knowledge, bidders' private information are now t_i, and independent, so revenue equivalence with (i) applies and profit is again $\left(\frac{n^2-n}{n^2+n}\right)$. The difference between (ii)b and (iii)b illustrates the linkage principle (see section 1.6).

(iv)(a) An equilibrium is for ascending auction behaviour to be unaffected. But it is not clear why bidders who know they do not have the highest value will bother to bid; if they don't, then profit will be much lower.

(b) The equilibrium is for the highest-value bidder to bid (and win) at (just above) the second-highest value. There is the same difficulty that low-value bidders may not bother to play.

(v) If the t_i are also affiliated we have

$\pi(\text{iva}) = \pi(\text{ivb})$ (assuming the equilibria described above),

$\pi(\text{iiia}) > \pi(\text{iiib})$ (affiliation),

$\pi(\text{iiib}) > \pi(\text{iib})$ (linkage), and

$\pi(\text{iva}) = \pi(\text{iiia}) = \pi(\text{iia})$ (trivially).

Summarising: $\pi(\text{iia}) = \pi(\text{iiia}) = \pi(\text{iva}) > \pi(\text{ivb}) > \pi(\text{iiib}) > \pi(\text{iib})$.

A2. (i) See Appendix 1.A.

The key is equation (1): $S_i(v) \geqslant S_i(\widetilde{v})+(v-\widetilde{v})P_i(\widetilde{v})$. This assumes private values. [It's not hard to do the common values case.]

(ii) [Equation numbers in the below all refer to Appendix 1.A; the Oxford examination candidates were not expected to write all of what follows!]

(a) (1) fails because a type v who deviates to behave as if he had type $\widetilde{v}$ does *not* earn the right hand side of (1) – because type v expects to face competitors with values that are conditional on her having type v, while type $\widetilde{v}$ expects to face competitors with values that are conditional on her having type $\widetilde{v}$. Appendix 1.C describes informally how to proceed, though candidates were not expected to go into this. A candidate might note that revenue equivalence still holds between the first-price sealed-bid and Dutch auctions, and also between the second-price sealed-bid and ascending auctions if either values are private or there are just two bidders.

(b) (1) fails because it is generally no longer true that

$$S_i(v) = vP_i(v) - i\text{'s expected payment.}$$

[One rather special case is that $S_i(v) = vP_i(v) - EU(v\text{'s payment})$

i.e., the bidder is risk-averse with respect to money, but the prize is not (equivalent to) money. In this case (1) holds as usual, hence $S(\cdot)$ is pinned down by $P(\cdot)$ and $S(\underline{v})$ as usual, so $EU(\text{type } v\text{'s payment})$ is equivalent across auctions under the usual revenue equivalence conditions. (And since risk-averse buyers are equally well off, a risk-neutral seller prefers the auction that stabilises their payments in order to give them a given expected utility most cheaply. So all-pay auctions raise more money than first-price auctions, which themselves raise more money than second-price auctions.)]

Again, of course, the first-price sealed-bid and Dutch are revenue equivalent, as are the second-price and ascending.

(c) This depends on the model of collusion – how do bidders share information and collude? But if, for example, all the bidders agree to (jointly) win the object at a price of zero, and then share the object among themselves à la Cramton, Gibbons and Klemperer (1987) then (1) holds and revenue equivalence follows as usual — it's just that each bidder's surplus is higher by the same amount $S(\underline{v}) > 0$.

A special case of CGK is McAfee and McMillan's (1992) suggestion that the n colluders allocate the object among themselves using a first-price "knockout" auction, the winner of which pays $1/n^{\text{th}}$ of his bid to all n colluders (including himself). This is incentive compatible (since a loser's payoff does not depend on his bid, each colluder makes his usual bid) and budget balanced ex post. Or the colluders can allocate the object among themselves in dominant strategies by running a second-price auction, the winner of which pays a risk-neutral ring-center who previously paid all the colluders $1/n^{\text{th}}$ of the expected second-highest of their n values (again a loser's payoff does not depend on his bid, so each colluder makes his usual bid), but this is only budget balanced in expectation – Graham and Marshall (1987). (The GM result extends to cases when not all bidders collude if the main auction is a second-price auction; if the knockout winner wins the main auction, he pays the ring-center any excess of the knockout price over the main auction price, and the ring-center previously pays all the colluders $1/n^{\text{th}}$ of the expected value of this.)

(Note, of course, that a colluder may want to cheat at the main auction. However, collusion is easier to sustain in second-price and ascending auctions than in first-price auctions, because in the former the designated winner can bid infinitely high and other colluders have no incentive to cheat and try to win the auction (Robinson, 1985). Also, (tacit) collusion is easier in ascending (multi-unit or repeated) auctions because bidders can use bids to signal (Chapters 3, 4, of my book; Klemperer, 2000; Brusco and Lopomo, 2002; Cramton and Schwartz, 2000). With non-independent values, an ascending auction allows colluding bidders to induce non-colluding opponents to bid less aggressively, by having some colluders drop out at a low price, thus signalling a low valuation (Pagnozzi, 2003). See also Section 1.9.)

(d) If each bidder is restricted to a single unit only, then equations (1) – (5), and hence revenue equivalence, hold as usual. The only change is that the $P_i(v)$ function is different to reflect the higher probability of winning that more units provides.

If a bidder's value is linear in the number of units won, we can just write $P_i(v) =$ (expected number of units i wins) and (1) – (5), and revenue equivalence, again hold as before.

With multiple identical indivisible units, and if bidders' values are not necessarily linear in the number won, a generalisation of the usual argument applies: let $\underline{v} \equiv (v_1, ..., v_N)$ represent the type whose *marginal* value of winning a j^{th} unit is v_j, and $\underline{P}(\underline{v}) \equiv (P_1(\underline{v}), ..., P_N(\underline{v}))$ be the vector of type $\underline{v}$'s equilibrium probabilities of winning at least j units (i.e., his value of winning exactly j units is $\sum_{k=1}^{j} v_k$, and his probability of doing so is $P_j(\underline{v}) - P_{j+1}(\underline{v})$). Then equation (1) becomes

$$S(\underline{v}) \geqslant S(\underline{\widetilde{v}}) + (\underline{v} - \underline{\widetilde{v}}) \cdot \underline{P}(\underline{\widetilde{v}}) \qquad (1')$$

(in which $(\underline{v} - \underline{\widetilde{v}}) \cdot \underline{P}(\underline{\widetilde{v}})$ is the dot product of $(\underline{v} - \underline{\widetilde{v}})$ and $P(\underline{\widetilde{v}})$) and the standard kind of argument now applies.[1]

Even more generally, if the units are not identical, let $\underline{v} \equiv (v_1, ..., v_L)$ be the type whose valuation of the j^{th} possible *allocation* of the units is v_j, and we again get equation $(1')$. (So with N distinct indivisible units and I bidders, $L = N^{I+1}$ if the auctioneer can retain units; a bidder's valuations may be different for allocations which differ only in the assignments to other bidders, so externalities between bidders can also be taken into account in this formulation.) (See Engelbrecht-Wiggans, 1988 and Krishna and Perry, 2000.)

So revenue equivalence applies fairly generally in the multi-unit context, in the sense that any two auctions that allocate the objects in the same way (i.e., have the same $\underline{P}(\underline{v})$ function) and give the same surplus to a particular type, are revenue-equivalent under the usual kinds of conditions.

However, it is important to note that standard auction forms will *not* in general result in the same allocation, $\underline{P}(\underline{v})$. In particular, while most standard auctions achieve the efficient allocation and are therefore revenue equivalent in the symmetric single-unit case, this is not true in the multi-object case because the demand-reduction and other problems both affect them and affect them differently (see Section 1.10 and the Afterword to Chapter 1).

[1] Substituting $\underline{v}$ by $t\underline{y}$ and $\underline{\widetilde{v}}$ by $(t+dt)\underline{y}$ into $(1')$, and then making the converse substitution, and taking $dt \to 0$, and integrating up in the usual way, yields

$$S(t\underline{y}) = S(\underline{0}) + \int_{\tau=0}^{t} \underline{y} \cdot \underline{P}(\tau\underline{y}) d\tau \qquad (5').$$

So the surplus function is pinned down by the $\underline{P}(\cdot)$ function, and we have revenue equivalence as usual.

For a simple example, consider 2 identical items for sale when each bidder has a two-dimensional type (v, k), values a single object at v, and values a second object at kv ($0 \leqslant k \leqslant 1$). Then
$S(v, k) \geqslant S(\tilde{v}, \tilde{k}) + (v - \tilde{v})P_1(\tilde{v}, \tilde{k}) + (kv - \tilde{k}\tilde{v})P_2(\tilde{v}, \tilde{k}) \quad (1')$

If we fix $k = \tilde{k}$ and write $\tilde{v} = v + dv$ then we can obtain an expression for $dv \to 0$:
$\frac{dS(v,k)}{dv} = P_1(v, k) + kP_2(v, k)$.

Similarly, we can fix $v = \tilde{v}$ and write $\tilde{k} = k + dk$ to obtain: $\frac{dS(v,k)}{dk} = vP_2(v, k)$.

These two equations are sufficient to pin down the surplus function, and hence ensure revenue equivalence across auctions that induce the same $P_1(\cdot, \cdot)$ and $P_2(\cdot, \cdot)$, for example, efficient auctions.

To illustrate, if $n = 2$, k is constant across bidders, v is distributed uniformly on [0,1], and the auction is efficient, the probability that bidder v wins at least 1 object is $\min(v/k, 1)$, and the probability that he wins both objects is kv. This gives $\frac{dS(v,k)}{dv} = \min(v/k, 1) + k^2 v$. For example, for $k = 0$, $\frac{dS}{dv} = 1$, because each bidder is guaranteed to win exactly one object; for $k = 1$, $\frac{dS}{dv} = 2v$, because $P_1 = P_2 = F(v) = v$.

(e) Equation (1) holds for types that exist with positive probability, but (4) (i.e., $\frac{dS(v)}{dv}$) may not be defined everywhere. However, the argument will go through, and so revenue equivalence will apply, unless the distribution is neither strictly increasing nor atomless, and fails only if there is an atom at the edge of a "gap":

When the distribution of bidders' feasible valuations has a gap from x^- to x^+ (i.e., the distribution is not strictly increasing in this range), the equation after (3) becomes: $P(x^+) \geqslant \frac{S(x^+)-S(x^-)}{x^+-x^-} \geqslant P(x^-)$

$$\Longrightarrow S(x^+) \in [S(x^-) + (x^+ - x^-)P(x^-), S(x^-) + (x^+ - x^-)P(x^+)]$$

When there is a gap but not an atom, $P(x^+) = P(x^-)$, so $S(x^+)$ is determined by $S(x^-)$ (and everywhere else it is determined by (4): $\frac{dS(v)}{dv} = P(v)$), so revenue equivalence holds.

Similarly, the surplus function is pinned down when there is an atom but not an adjacent gap. Suppose there is an atom at x. Then there will be a discontinuity in the $P(\cdot)$ function at x: bidders with values slightly higher than x win against the positive measure set of bidders with value x, whereas bidders with values slightly lower than x do not. When there is not also an adjacent gap, all this means is that there is a discontinuity in $\frac{dS(v)}{dv}$ and hence a kink in the $S(\cdot)$ function at x, so the entire surplus function is pinned down, and revenue equivalence again holds.

When there is an atom at the edge of a gap, however, it is impossible to pin down the surplus function: $P(x^-) < P(x^+)$, so $S(x^+)$ is not determined by $S(x^-)$, and revenue equivalence fails. See (the solution to) Exercise 2 of *Auctions: Theory and Practice*, which we discuss below.[2]

Nevertheless the standard auction forms generally remain revenue equivalent even when there are atoms at the edges of gaps (and hence also mixed strategy equilibria in, for example, first-price auctions).

For example, consider the case of N symmetric bidders each of which has a private value independently drawn from a finite set of types, $v_1 < .... < v_H$ with probabilities $p_1, ..., p_H$ respectively. Let $P(v_k) = \left(\sum_{j=1}^{k-1} p_j\right)^{N-1}$ (so $P(v_k)$ is the probability that a bidder with value v_k has a strictly higher value than all his $N-1$ competitors). In a first-price auction, type v_1 bids v_1, while other types

[2]A simple illustration is a single bidder whose value is L or H, to whom the seller sells with probabilities $P(L) = 0$ and $P(H) = 1$, respectively, by offering price $r \in (L, H]$. Then $S(L) = 0$, and $S(H) \in [0, H - L]$ and the left and right ends of the interval are achieved by $r = H$ and $r = L$ respectively.

randomise. By standard arguments there can be no atoms in the randomisation, nor can there be gaps (since the lowest type above a gap would lower his bid), nor can the ranges of different types overlap (since an overlap would imply that a higher and a lower type of the same bidder are both happy choosing these different strategies that give the higher type a lower probability of winning[3]). So type v_k is indifferent about mimicking type v_{k-1} by bidding at the top of type v_{k-1}'s range (which is the bottom of type v_k's range). But in an ascending auction, also, type v_k is indifferent about (almost) mimicking type v_{k-1} by dropping out at (just above) v_{k-1}. So in both these auctions

$$\begin{aligned} S(v_k) &= S(v_{k-1}) + (v_k - v_{k-1})P(v_k) \\ \Longrightarrow S(v_k) &= S(v_1) + \textstyle\sum_{j=1}^{k}(v_j - v_{j-1})P(v_j) \qquad (5') \end{aligned}$$

That is, bidder surplus is pinned down and equal across the standard auctions in the "discrete" (finite number of types) case, just as in the "continuous" case (in which the distribution from which bidders' types are drawn is strictly increasing).

The point is that these auctions have the property that the highest type always wins so (5′) approximates (5) in a continuous approximation to the discrete case. Although we can find examples of auctions in the discrete case that are efficient but not revenue equivalent to the standard auctions (and may raise more money than the standard auctions), these examples don't yield efficiency in a continuous approximation to the discrete model that "fills in the gaps".

To illustrate all these issues, consider Exercise 2's analysis of efficient mechanisms which always assign a unit to just one of two bidders who each have value v_H with probability p_H and value v_L with probability p_L. Efficiency implies types v_H and v_L win with probabilities $P(v_H) = \left(\frac{p_H}{2} + p_L\right)$ and $P(v_L) = \frac{P_L}{2}$ respectively, and (1) $\Longrightarrow$

$$\begin{aligned} S(v_H) &\geqslant S(v_L) + (v_H - v_L)P(v_L) \qquad (1^*) \\ \text{and } S(v_H) &\leqslant S(v_L) + (v_H - v_L)P(v_H) \qquad (1^{**}) \end{aligned}$$

giving a range of solutions from the seller-optimum characterised by (1*) holding with equality, to the buyer-optimum characterised by (1**) holding with equality. (See solution to Exercise 2.)

Now consider a continuous approximation of this discrete case in which most types are very close to v_H or v_L with a tiny density everywhere else. Revenue equivalence holds, of course, for this distribution for any given allocation. Furthermore, for the types v_H^- (the lowest of the types who are very close to v_H) and v_L^+ (the highest of the types who are very close to v_L), for example

$$S(v_H^-) \geqslant S(v_L^+) + P(v_L^+)(v_H^- - v_L^+) \qquad (1^{\tilde{}})$$

[3]To see this is a contradiction, see the equation below (3) (which does not require small dv) in Appendix 1A.

$$\text{and } S(v_H^-) \leqslant S(v_L^+) + P(v_H^-)(v_H^- - v_L^+) \qquad (1^{\sim\sim})$$

Of course, $S(v_H^-) \approx S(v_H)$ and $S(v_L^+) \approx S(v_L)$ (using equation (1), since $v_H - v_H^- \approx 0$ and $v_L^+ - v_L \approx 0$). But pinning down the allocation in the discrete case (all v_H's beat all v_L's) does *not* pin down a unique allocation in the continuous analog: an auction can treat types such as v_H^- and v_L^+ in several different ways while still corresponding to the (symmetric) efficient allocation of the discrete case.

To find the seller-optimum in the continuous case, note that the marginal revenue is initially just below v_H, then turns very negative at v_H^- when the density suddenly becomes very thin, then becomes barely positive again at v_L^+. So to maximise revenue, all types with a value of v_H^- or below are pooled and given an equal chance of winning.[4] (See Bulow and Roberts (1989); as always, we cannot give higher-value types lower chances of winning than lower-value types.) So $P(v_H^-) = P(v_L^+) = P(v_L)$. (Of course, $P(v_H)$ and $P(v_L)$ are unchanged from the discrete case.) So $(1^{\sim})$ and $(1^{\sim\sim})$ both hold with equality, and therefore $S(v_H) = S(v_L) + (v_H - v_L)P(v_L)$, just as in the seller-optimum of the discrete case.

The "seller-pessimum" of the continuous model, by contrast, minimises revenue by pooling all the negative marginal revenues in with the high marginal revenue bidders at the top of the distribution, i.e., pools everyone at v_L^+ or above. (Again, we cannot give higher-value types lower chances of winning than lower-value types.) So now $P(v_L^+) = P(v_H^-) = P(v_H)$ and $(1^{\sim})$ and $(1^{\sim\sim})$ again hold with equality, but in this case $S(v_H) = S(v_L) + (v_H - v_L)P(v_H)$ as in the buyer optimum of the discrete model.

Finally consider the continuous model for any standard auction (e.g., ascending or first-price sealed-bid) in which the highest-value bidder wins. In this case $P(v_H^-) \approx P(v_L^+) = p_L > P(v_L)$ (recall $P(v_L) = \frac{p_L}{2}$) since the intermediate types beat all v_L's whereas a v_L beats another v_L only half the time. So $(1^{\sim})$ and $(1^{\sim\sim})$ both hold with (approximate) equality and $S(v_H) = S(v_L) + (v_H - v_L)p_L$. Again this corresponds exactly to a standard auction such as a first-price or ascending auction in the discrete model in which, as noted above: (a) a type v_H is just indifferent about (almost) mimicking a v_L, and (b) a type v_H who (almost) mimicks a v_L beats all v_L's and no v_H's so wins with probability p_L, so again $S(v_H) = S(v_L) + (v_H - v_L)p_L$.

[4] In the optimal auction the price starts at the pooling price of v_L, and then jumps to $\approx \frac{v_H + v_L}{2}$. (This makes v_H equally well off as if he joins the pool by bidding v_L and then wins with probability 1/2 when the other bidder bids v_L. He (almost) always makes no money if the other bidder is not a v_L.) When two bidders jump, the price rises continuously to (usually) about v_H.

This optimal auction can alternatively be implemented by first offering a price, p, that makes v_H just indifferent between pooling at v_L and accepting (and possibly being rationed at) p, i.e., $\left(\frac{p_H}{2} + p_L\right)(v_H - p) = \frac{p_L}{2}(v_H - v_L)$; and then offering price v_L if no-one accepts p. This is the form in which we described the mechanism in the solution to Exercise 2. Strictly, this alternative implementation is slightly suboptimal for an atomless distributions of types. But it is arbitrarily close to optimal for distributions of types that arbitrarily closely approximate the "discrete" two-type distribution. And it *is* optimal if (as suffices for our purposes) we restrict attention to strictly increasing distributions of types that put atoms at v_H and v_L and a small density everywhere else.

The way I think of all this, then, is that revenue equivalence holds in the discrete case over mechanisms that would induce the same allocation across all types including those types that do not actually exist (but would exist if we filled the gaps with a continuous approximation): if a type v_H^- were to exist in our discrete example he would pool with the L's in the seller-optimum; if a type v_L^+ existed he would pool with the H's in the seller-pessimum; but either of these types would beat the L's and lose to the H's in a "standard" auction.

References

Brusco, S. and Lopomo, G. (2002) Collusion via Signalling in Simultaneous Ascending Bid Auctions with Heterogeneous Objects, With or Without Complementarities. *Review of Economic Studies*, 69, 407-436.

Bulow, J. and Roberts, J. (1989) The Simple Economics of Optimal Auctions, *Journal of Political Economy*, 97, 1060-1090.

Cramton, P., Gibbons, R., and Klemperer, P. D. (1987) Dissolving a Partnership Efficiently, *Econometrica*, 55, 615-632.

Cramton, P. and Schwartz, J. A. (2000) Collusive Bidding: Lessons from the FCC Spectrum Auctions. *Journal of Regulatory Economics*, 17, 229-252.

Engelbrecht-Wiggans, R. (1988) Revenue Equivalence in Multi-Object Auctions. *Economics Letters*, 26, 15-19.

Graham, D. A and Marshall, R. C. (1987) Collusive Bidder Behavior at Single-Object Second-Price and English Auctions. *Journal of Political Economy*, 95, 1217-1239.

Klemperer, P. D. (2000) What Really Matters in Auction Design. Working paper, Nuffield College, Oxford (revised version reprinted as Chapter 3 of *Auctions: Theory and Practice*).

Krishna, V. and Perry, M. (2000) Efficient Mechanism Design. Working paper, Pennsylvania State University and Hebrew University of Jerusalem.

Maskin, E. S. and Riley, J. G. (1985) Auction Theory with Private Values. *American Economic Review*, 75, 150-155.

McAfee, R. P. and McMillan, J. (1992) Bidding Rings. *American Economic Review*, 82, 579-599.

Pagnozzi, M. (2003) Bids as a Vehicle of (Mis)Information: Collusion in English Auctions with Affiliated Values. Mimeo, Oxford University.

Robinson, M. S. (1985) Collusion and the Choice of Auction. *The RAND Journal of Economics*, 16, 141-145.

A3. [A closely related question is Problem 9. See also Problem 1 (and Part 3 of my first-year course notes).]

(i) By symmetry i will win if $z_i \geqslant z_j$.

So i will quit at $p = z_i + z_i = 2z_i$.

[i is just indifferent about quitting or not at this point. It is easy to see that if i were to win later she would lose money (assuming j is following her equilibrium strategy) and quitting earlier would mean quitting when i's value exceeds the asking price (again assuming j is following her equilibrium strategy). See Appendix 1.D (or section 3.2 of my first-year course notes).]

(ii) If i has type z, then conditional on i winning, j is uniformly distributed as $[0, z]$. So on average j will quit at $2\left(\frac{z}{2}\right)$ conditional on i winning.

i.e. i's expected payments are z conditional on winning.

By the Revenue Equivalence Theorem, the expected payment of type z is the same in the English and Dutch auctions, so it is the same conditional on winning (since type z has the same probability of winning either auction).

So type z bids z in the Dutch auction.

(iii) Type z's expected payments are z^2 unconditional on winning in the English (and Dutch) auction.

By the Revenue Equivalence Theorem, the expected payment of type z is the same in the English (and Dutch) and All-Pay auctions, so type z bids z^2 in the All-Pay auction.

(iv) We look for a symmetric equilibrium in which type z bids $b(z)$.

If z bids $b(\widetilde{z})$, she beats types $\leqslant \widetilde{z}$, i.e. with average signal $\frac{\widetilde{z}}{2}$. So conditional on winning type z's value is $z + \frac{\widetilde{z}}{2}$, and she wins with probability $\widetilde{z}$, so type z's expected surplus from bidding $b(\widetilde{z})$ is $S = (z + \frac{\widetilde{z}}{2})\widetilde{z} - b(\widetilde{z})$.

$\frac{\partial S}{\partial \widetilde{z}} = z + \frac{\widetilde{z}}{2} + \frac{\widetilde{z}}{2} - b'(\widetilde{z})$

$= 0$ at $\widetilde{z} = z$ in equilibrium

$\Longrightarrow b'(z) = 2z$ with boundary condition $b(0) = 0$ (since $z = 0$ never wins in equilibrium). The solution is $b(z) = z^2$ which agrees with part (iii).

(v) See Klemperer (1998) (or section 3.5 of my first-year course notes): bidder 2 quits at z_2 in the ascending auction (if she bothers to enter at all); the Dutch auction solution is almost unaffected from part (ii).

For revenue equivalence between two auctions to hold, we need that the probability that any given type of bidder will win is the same in the two auctions.

See Appendix 1.A. (It would be sufficient for the purposes here that the object always went to the bidder with the highest signal.) But here type z of bidder 1 wins an ascending auction with probability 1, but wins a Dutch auction with probability only (slightly above) z, see section 1.7.2 of the book or Klemperer (EER 1998) (or section 3.5 of my first-year course notes).

A4. (i)(a) Initially each bidder stays in until it would be just indifferent if it suddenly found itself a winner, i.e., if two others quit simultaneously. So at the point at which a bidder with value t_i quits the price must be

$$p = t_i + [2t_i + E(t \mid t \geq t_i)] \tag{1}$$

(i.e., if two others were to quit, their signals would be assumed to be t_i and the remaining signal would be assumed to be somewhere above t_i). So after the first quit at price p, the signal of the quitter is inferred to be the t_i that satisfies (1). Call this signal $t_{(4)}$. Remaining bidders with signals t_i then quit at prices such that they would be just indifferent about finding themselves winners, i.e., at prices

$$p = t_i + \left[t_{(4)} + t_i + E(t \mid t \geq t_i)\right] \tag{2}$$

(i)(b) If a bidder with value t_i were suddenly to find himself a winner at the point at which he was himself about to quit he would infer the quitter's signal would also be t_i, and the two bidders excluded from his auction had signals $E(t)$. i.e. he quits at price

$$p = t_i + [t_i + 2E(t)] \tag{3}$$

(i)(c) In the single-auction case, (a), expected revenue $= 2E\{t_{(3)} + \left[t_{(4)} + t_{(3)} + E(t \mid t \geq t_{(3)})\right]\}$ using (2). For the uniform distribution, $Et_{(i)} = \frac{5-i}{5}$, so expected revenue $= 2(\frac{2}{5} + \frac{1}{5} + \frac{2}{5} + \frac{3\frac{1}{2}}{5}) = 3\frac{2}{5}$. And in the two-auction case, (b), expected revenue $= 2E\{t_{(2 \text{ of } 2)} + t_{(2 \text{ of } 2)} + 2E(t)\}$ in which $t_{(2 \text{ of } 2)}$ is the second-highest of two signals, using (3), so for the uniform distribution expected revenue $= 2(\frac{1}{3} + \left[\frac{1}{3} + 2(\frac{1}{2})\right]) = 3\frac{1}{3}$.

(ii) (a) In the single-auction case the expected revenue is $E(MR_{(1)} + MR_{(2)})$. In the two auction case the two high-signal bidders are in the same auction one-third of the time (in which case the winner are the highest and third-highest signal bidder), so the expected revenue is $E(MR_{(1)} + \frac{2}{3}MR_{(2)} + \frac{1}{3}MR_{(3)})$. So revenue is higher in the single-auction case if $E(MR_{(2)}) > E(MR_{(3)})$.

(b) (I) For $F(t) = t, MR(t) = v + t - 1$, so $MR_{(2)} > MR_{(3)}$ always, so the single-auction case is more profitable.

(II) For $F(t) = 1 - t^{-2}$, $MR(t) = v - \frac{t}{2}$, so $MR_{(2)} < MR_{(3)}$ always, so the two-auction sales mechanism is more profitable.[5]

(III) Since the problem is pure common values, any allocation is equally efficient, so bidders' preferences are precisely the opposite of the seller's.

(IV) In the private-value case, $v_i = t_i$, marginal revenue $MR_i = t_i - \frac{1-F(t_i)}{f(t_i)}$ is "downward sloping" for both $F(t) = t$ $(MR(t) = 2t - 1)$ and $F(t) = 1 - t^{-2}$ $(MR(t) = \frac{t}{2})$, so the seller prefers the single auction for either distribution.

(c) A good exam answer would just be that non-decreasing MR is much more likely with common than with private values, so it's much more likely that separate auctions are more profitable with common than with private values.

A good answer might also point out why: with pure common values (our problem) non-decreasing MR just means $(-\frac{1-F(t_i)}{f(t_i)})$ is non-*in*creasing in t_i which is obviously much easier than having $(t_i - \frac{1-F(t_i)}{f(t_i)})$ non-*in*creasing. We could note that non-decreasing MR is even easier in common-value contexts if values are non-additive (e.g., we could discuss the "maximum game" see Bulow and Klemperer "Prices and the Winner's Curse", *Rand Journal*, 2002). With private values, by contrast, we know $t_{(2)} = E(MR(t) \mid t \geq t_{(2)})$ so $E(t_{(2)}) = E(MR_{(1)})$ and, of course, $t_{(2)} > MR_{(2)}$, so $E(MR_{(1)}) > E(MR_{(2)})$ which is at least suggestive of the difficulty of having sufficiently non-decreasing MR for split auctions to be more profitable (but see below).

I did not expect any examinee to write any of what follows:

[5] Note that MR's are the easier way of computing expected revenues. For the uniform case and a single-auction,

$$\begin{aligned} \text{expected revenue} &= E(MR_{(1)} + MR_{(2)}) \\ &= E([v + t_{(1)} - 1] + [v + t_{(2)} - 1]) \\ &= ([2 + \frac{4}{5} - 1] + [2 + \frac{2}{5} - 1]) = 3\frac{2}{5} \end{aligned}$$

while with two auctions, expected revenue is reduced by $\frac{1}{3}E(MR_{(2)} - MR_{(3)})$, i.e., by $\frac{1}{15}$ - we obtain the result without needing to compute exactly where bidders quit the auction.

For $F(t) = 1 - t^{-2}$, it is equally straightforward to compute profits using marginal revenues but considerably more tiresome (for me at least) to compute prices directly as in part (i).

One might even wonder whether separate auctions could ever be more profitable in the private case since always $E(MR_{(1)}) > E(MR_{(2)})$. However, it is *not* always true that $E(MR_{(2)}) > E(MR_{(3)})$, because $t_{(3)} = E(MR(t) \mid t \geq t_{(3)})$, so $E(t_{(3)}) = \frac{1}{2}E(MR_{(1)} + MR_{(2)})$ (since $t_{(3)}$ is the price, i.e. average revenue, obtained from selling to everyone above $t_{(3)}$ who on average are $t_{(1)}$ and $t_{(2)}$), so $E(MR_{(2)}) = E(2t_{(3)} - t_{(2)}) \neq E(t_{(3)})$.

A simple example which give separate auctions being more profitable with private values is a two-point support: probability 10% of α; probability 90% of $\beta < \alpha$, so Prob $(t_{(4)} = \alpha) = 0.0001$, Prob $(t_{(3)} = \alpha) = 0.0037$, Prob $(t_{(2)} = \alpha) = 0.0523$, etc.[6] For this private-values example, it's actually easier to note directly that expected revenue from a single auction is $E(2t_{(3)})$ while that from two separate auctions is $E(\frac{1}{3}(t_{(2)}+t_{(4)}+\frac{2}{3}(t_{(3)}+t_{(4)}))$ so the separate auctions win iff $E(t_{(2)}-t_{(3)}) > 3E(t_{(3)}-t_{(4)})$. (Of course it's easy to check that the MR formula yields this too.[7])

Returning to private values *vs* common values: since $E(t \mid t \geq t_{(3)}) = E\left(\frac{t_{(1)}+t_{(2)}}{2}\right)$ while $E(t_{(2\text{ of }2)}) = E(\frac{1}{6}(t_{(2)}+t_{(4)})+\frac{1}{3}(t_{(3)}+t_{(4)}))$ we find (from part (i)) that two separate auctions are more profitable in our common value case iff $2E(t_{(2)}-t_{(3)}) > 3E(t_{(3)}-t_{(4)})$ which is clearly easier to satisfy than the private-value condition.

It's also easy to look at bidders' preferences in the private-value case. Bidder surplus is welfare less revenue, i.e., $v - MR$ of the winning bidders (in expectation). So the bidders prefer the single auction iff

$$E(t_{(2)} - MR_{(2)}) > E(t_{(3)} - MR_{(3)})$$

or $E\left(\frac{1-F(t_{(2)})}{f(t_{(2)})}\right) > E\left(\frac{1-F(t_{(3)})}{f(t_{(3)})}\right)$. Note that this is *precisely* the condition for the seller to prefer the two auctions in our additive common-value model. With linear demand this is false (because MR is twice as steep as demand everywhere) but with CE -2 demand ($F(t) = 1 - t^{-2}$) it is true (because MR is half as steep as demand everywhere). With CE -2 demand, bidders with private values dislike the split auction because the efficiency losses outweigh the small reduction in payments to the seller.

Those who are truly fascinated by this problem may wish to look at Ellison, Fudenberg and Möbius' paper "Competing Auctions", *Journal of the European Economic Association* (2004) though this paper is somewhat differently focused. (Möbius is a former Oxford Economics M.Phil student who is now (2006) a Harvard Professor).

[6] Another example is inelastic constant-elasticity demand. If $F(t) = 1 - t^{\eta}$ then $MR(t) = t(1 + \frac{1}{\eta})$ which is negative and therefore increasing everywhere (as t falls) if $\eta = -\frac{1}{2}$, say. This is rather confusing, especially since I've just pointed out that $E(MR_{(1)}) > E(MR_{(2)})$. The problem is that marginal revenue is really infinite at $t = 1$ in this case, reflecting the infinite expected profits that can be made by selling with arbitrarily low probability. This can all be made clear and sensible (i.e. expected profits always finite) by considering a close approximation that hits the price axis at a finite price.

[7] $E(t_{(2)}) = E(MR_{(1)}); E(t_{(3)}) = E\left(\frac{MR_{(1)}+MR_{(2)}}{2}\right)$;
$E(t_{(4)}) = E\left(\frac{MR_{(1)}+MR_{(2)}+MR_{(3)}}{3}\right)$.

A5. SKETCH ANSWER

The bare bones of an answer are discussed in the book (*Auctions: Theory and Practice*, Klemperer (2004)) see e.g., pages 33 (including note 91) and 63-4 (including note 7). Other useful references are Rothkopf et al (JPE 1990), Milgrom's 2005 *Clarendon Lectures* (given in Oxford this academic year), and Ausubel and Milgrom "The Lovely but Lonely Vickrey Auction" (Chapter 2 of *Combinatorial Auctions*, Cramton, Shoham and Steinberg,eds.(2006)).

The efficiency properties in the private-values context are standard.

Problems that apply even in the single-unit context include
- others can see the winner's value (so the winner may be vulnerable to ex-post exploitation (Rothkopf, JPE, 1990), or auctioneer exploitation (the auctioneer may have a confederate place a bid just less good than the winner's to leave the winner (almost) zero surplus), or even absent foul play the result may be a PR disaster (as in the celebrated New Zealand auction, pg 110 of the book)
- the efficiency of the auction can facilitate entry deterrence by strong players and be disastrous for revenue (see the general discussion of entry in 1st vs 2nd price auctions in section 4.3.1 of the book and elsewhere)
- disadvantages of 2nd price auctions vs 1st price auctions more generally can be discussed (revenue properties with risk-aversion, asymmetries, budget constraints, etc; see book)
- efficiency properties are limited outside the private-value context.

In the multi-unit context there are more problems including:
- the auction is "non-transparent", that is, payments are hard to understand for non-experts, as is the mechanism itself for many people
- different bidders pay different amounts for identical objects and often those who value an object more pay less
- in multi-unit contexts the mechanism "favours larger bidders over smaller bidders" (e.g. relative to a standard ascending auction)
- when there are complementarities, there are curious and difficult-to-guard-against collusive possibilities (even for bidders who would all lose absent collusion)
- collusion may be facilitated even absent complementarities (though ascending auctions are worse for collusion -see the book)
- demerging, or setting up new bidders, can help a bidder
- revenues can be very low, when goods are complements (and, relatedly, adding bidders can sometimes *reduce* revenues) as discussed in, for example, Milgrom's Clarendon lectures.

A6. [I have ignored subscripts.]

(i) Total bidder surplus is the expected high value of 2/3, since expected transfers have to sum to zero. So surplus per bidder is on average 1/3 versus 1/6 in a "regular" auction.

(ii) The revenue equivalence theorem applies exactly, and for the same reasons, as "normal". So

$$S(v) = S(0) + \int_0^v P(x)dx$$

Clearly, $P(v) = v$. Also, given the symmetry of the problem, $S(0) = 1/6$ because average surplus is 1/6 greater than in the "regular" auction (1/3 versus 1/6) and we know from the revenue equivalence theorem that the difference is a constant type-by-type since $dS/dv = P(v)$ (the same amount) in both games, for all types. So

$$S(v) = S(0) + \frac{v^2}{2} = \frac{1}{6} + \frac{v^2}{2}.$$

(If necessary, candidates could work out that $S(0) = \frac{1}{6}$, by using the fact that $E(S(v)) = \frac{1}{3}$.)

(iii) In the sealed-bid auction the bidder will pay his bid when he is highest and receive the competitor's bid when that bid is higher. For a bidder with value v we have

$$1/6 + v^2/2 = S(v) = v^2 - vb(v) + \int_v^1 b(x)dx$$

since the right-hand side computes expected surplus as value times probability of buying, less expected payments when a buyer, plus expected receipts when a seller. We can now see directly which linear function solves this equation, or we can differentiate both sides to obtain

$$v = 2v - b(v) - v\frac{\partial b}{\partial v} - b(v)$$

which yields $\frac{\partial b}{\partial v} = 1 - 2\frac{b(v)}{v}$. Since the question makes clear that we will find $\frac{\partial b}{\partial v}$ is a constant it is easy to see that a solution is $b(v) = \frac{v}{3}$. (It is also clear directly from the formula for $S(v)$ that $b(1) = 1/3$, since a bidder with value 1 must be a winner.)

(iv) In the ascending auction the analogous surplus equation is:

$$1/6 + v^2/2 = S(v) = v^2 - \int_0^v b(x)dx + (1-v)b(v)$$

which produces (through differentiation of both sides)

$$v = 2v - 2b(v) + (1-v)\frac{\partial b}{\partial v}$$

so $\frac{\partial b}{\partial v} = \frac{2b(v)-v}{1-v}$. The formula immediately yields $b(0) = 1/6$, and the solution is therefore $b(v) = \frac{1}{6} + \frac{v}{3}$.

Bidders go beyond their value, and take the risk of buying at a price that exceeds their value, in the hope of selling out at a higher value (as in Bulow, Huang and Klemperer (JPE 1999) though in a private value context).

(v) Type v of a partner with share α earns $1/6+v^2/2-\alpha v$, so it is simple to compute that the worst off type of a partner with share α has value α and earns $1/6-\alpha^2/2$ (or $1/24$ for $\alpha = 1/2$ — up to now we are assuming $\alpha = 1/2$, but we will want general α in the next part.)

(vi) Nothing is different in the analysis of general shares. If we are to run an efficient mechanism, so $P(v) = v$, incentive compatibility requires us to maintain $S(v) = S(0) + \frac{v^2}{2}$. All we can do is to alter $S(0)$ by having the smaller partner pay the larger partner for participating in the mechanism.

Given shares α and $(1-\alpha)$, the worst-off types of the two partners earn surpluses of $1/6-\alpha^2/2$ and $1/6-(1-\alpha)^2/2$ respectively. Therefore, we can have the smaller partner pay $\left|\frac{1}{4}-\frac{\alpha}{2}\right|$ (i.e., half the difference between these two surpluses) to the larger partner, so that both worst-off types make (the same) non-negative surplus if $(1/6-\alpha^2/2) + (1/6-(1-\alpha)^2/2) \geq 0$ — i.e., $\alpha \in [1/2-\sqrt{3}/6,\ 1/2+\sqrt{3}/6]$. For α outside this interval, efficient individually-rational dissolution is impossible if the partners know their own values before agreeing to a mechanism. (This confirms Myerson and Satterthwaite's (JET 1983) result on the impossibility of efficient bilateral trading when one player owns 100% of the good to be traded.)

See Cramton, Gibbons, and Klemperer (Econometrica 1987) for the general case; note also that the problem is identical to one in which the two bidders collude in buying from a seller, by bidding against each other in a pre-auction knock-out for the right to buy the asset from the seller for zero.

A7. [Sketch Answer]

(i) [Paul Klemperer was quoted in the Economist 9/9/06, p77.] He was probably suggesting that auctioning larger blocks of spectrum might have worked better in this particular case.

(ii) concern about "exposure" among some bidders. That is, bidders who had a strong preference for complete coverage of a region were reluctant to bid on smaller blocks for fear that they would be "stranded" winning some but not all parts of the region. The result could be that bidders who either wanted only small blocks, or were willing to risk winning only some parts of the region could have picked up those areas cheaply. The fear of stranding would be reinforced if resale was not permitted, or was hard for some other reason. Larger bidding increments are perhaps slightly more likely to create this problem.

(iii) a Vickrey auction would sell private-value objects efficiently at disparate prices assuming a fixed number of non-colluding bidders, private values, and no budget constraints.

A sequential auction (by English/Japanese, Dutch, first-price sealed bid, second-price sealed bid, etc) or a discriminatory auction of identical objects is, in theory, efficient assuming symmetric bidders.

Perfect price discrimination is efficient, and imperfect price discrimination may be more efficient than a fixed price.

(iv) if there are no complementarities at all between the smaller areas, then the objects should be sold independently, but if the smaller areas are perfect complements (individually useless to an owner without the whole region) then they should be sold together. More generally, selling a small number of large licences makes life very difficult for bidders who only want smaller areas, especially if resale is hard. So it depends upon whether it is more efficient to encourage bidders who are interested in winning smaller areas (these bidders may be new entrants) or bidders interested in winning larger regions.

(v) in a package auction, a bidder can specify a bid for a group of properties and then wins all of these if his bid beats the sum of the best bids for the individual properties. This still makes life very difficult for smaller bidders who have difficulty coordinating their bids to beat the package bidder, and also face an "free-rider" problem (each small bidder wants the other small bidders to bid higher prices that will contribute to defeating the package bidder). There are other problems: a bidder who wants to bid for a package as well as for individual units is bidding against himself; if bids must be left on the table in subsequent rounds then bidders may be reluctant to make offers, but if bids need not be left on the table coordination is particularly hard, etc.

A8. [Sketch of a possible answer.]

1. The candidate might begin by showing that under the appropriate conditions, all auctions make the same expected profits

(see sec 2 of my first-year class notes on the course website)

2. But the profits are only the same in expectation,

and under the assumptions that:

- The number of bidders is fixed
- They play noncooperatively
- There is a single indivisible unit (or if there are multiple identical objects, bidders want a single indivisible unit each)
- Bidders' values are independent
- They are risk neutral
- There are no budget constraints
- There are no externalities between bidders
- Bidders are drawn from symmetric distributions, and the auction is a standard one, so that the highest actual type is the winner (or some alternative assumptions so that the the lowest type's payoff is constant (e.g., zero), and a given type of a given bidder wins any auction with equal probability).

Note that setting a reserve price affects types' probabilities of winning – so setting a reserve price is an important part of auction design, even if all the other assumptions above hold.

(Not all these issues were discussed in the first-year lectures, but many were)

3. The candidate might then discuss how auctions' expected profits differ if these assumptions are relaxed. (With the exception of independence, probably externalities, and in some cases symmetry, relaxing the assumptions above generally favours first price over second price mechanisms.)

(only some of this was discussed in the first-year lectures)

4. The auctioneer's objective may not be (only) revenue. (In particular, efficiency considerations tend to favour second price mechanisms.)

A9. (i) See Appendix 1.C of *Auctions: Theory and Practice*. See also Section 1.6.

(ii) References, and the bare bones of an answer are in Section 1.6.

(iii) See Section 1.6, Appendix 1.C and also, e.g., Section 4.2.

See also question 17 of the book and its sketch solution.

A10. (i) A. We run a "second marginal-revenue auction". The woman's demand curve is $80-40q$, so her marginal revenue is $80-80q = 80-80((80-v)/40) = 2v-80$, while the man's demand is $40-40q$, so his $MR = 2v-40$. So we ask the bidders to state their values, and give the object to the bidder with the higher marginal revenue, provided that marginal revenue exceeds the seller's valuation, at a price equal to the lowest value at which that bidder could have won given the other bidder's report (and given the seller's cost).

B. Since the woman's marginal revenue beats the man if and only if her value exceeds his by at least 20, we could give the man a discount of 20 conditional on him being the winner, and simply run a standard ascending auction with a reserve price of 45.

C. An optimal auction maximises expected marginal revenue, just as a price discriminating monopolist maximises the sum of marginal revenues. Because the man is drawn from a distribution that is lower by a constant, he has a higher marginal revenue at any given value.

(ii) A. This is the problem described in Bulow and Roberts (1989), pages 1078-81, but with 20 deducted from all the valuations. The demand curve for $q < 20$ is $80-200q$, so $MR = 80-400q = 2v-80$ (as before); demand for $q > 20$ is $50-50q$, so $MR = 50-100q = 2v-50$. Because marginal revenue is not downward sloping, selling to the highest value bidder does not sell to the highest marginal-revenue bidder, so the seller would do better to "iron" the marginal revenue curve and can do this best by treating all bidders with values between 35 and 50 equally, as if they all had marginal revenues of 20 (since the triangular areas cut off by the $MR = 20$ line are then equal in a diagram like Bulow and Roberts' figure 3). The reserve price is 30, since a bidder with value 30 has $MR = 10$.

The higher-value bidder is sold to at the price of the lower-value bidder except that the randomisation (if it is necessary) is at a price of 35, and that if the winner's value exceeds 50 and the runner-up's value is between 35 and 50, the winner pays 42.5 (so that if the winner had a value of 50, as was just necessary to guarantee victory, it would have been indifferent between its outright victory and participating in the lottery).

[The question did not ask this, but this outcome could be achieved by running a simple ascending auction, but skipping all the prices between 35 and 42.5. In fact, no-one will then quit between 35 and 50; any bidder with a value between 35 and 50 will quit at 35–if its competitor quits simultaneously, it prefers a fair lottery at 35 to winning for sure at 42.5, and if its competitor doesn't quit simultaneously, it is doomed anyway. (With n bidders still in the auction at a price of 35, the auctioneer would need to jump the price to $50 - (15/n)$.)]

B. It doesn't matter whether or not we treat the man and woman fairly, provided that for each individual we make it equally likely that all values between 35 and 50 are equally likely to win, so that the expected MR of the winner is as in the case of "fair" allocation (and we must choose prices so that this outcome is achieved given the incentive compatibility constraints).

If we always favour the man in the "ironing" range, then the higher-value bidder is sold to at the price of the lower-value bidder (as before), except that: if both values are between 35 and 50 the man wins at a price of 35; if the man's value exceeds 50 and the woman's value is between 35 and 50, the man pays 35 (so if the man had a value of 50, he would have been indifferent between stating that value and stating a somewhat lower value); and if the woman's value exceeds 50 and the man's value is between 35 and 50, the woman pays 50 (so if the woman had a value of 50, she would have been indifferent between stating that value and stating a somewhat lower value).

A11. (i) see part 2 of the course notes on the web (or Appendix 1A of my book).

(iia) $V/2$ (= expected second highest signal).
profit $= V/4$ (= expected difference between highest and second highest signal).

[Recall that the expected value of the k^{th} highest of n independent draws from a uniform distribution on $[X, Y]$ equals $X+[(n+1-k)/(n+1)](Y-X)$ from the course notes on the web, or Appendix 1A of my book, or this is not hard to compute directly, for $n = 3$ firms.]

(iib) $V/2$, and profit $= V/4$ by revenue equivalence theorem with ascending auction.

[Recall revenue equivalence applies to expected bidders' surpluses as well as to expected total payments.]

(iiia) the contest for the uncommitted consumers is a sealed bid auction which yields $(\theta - \varphi)C/4$, by revenue equivalence with ascending auction. All three firms also make $\varphi(1 - c_i)$ automatically on their locked in customers, and average costs are $C/2$, so expected industry profits are $3\varphi(1 - C/2) + (\theta - \varphi)C/4$.

(iiib) prize is $(\theta - \varphi)(1 - c_i)$, independent – of course – of the bids. Payment is $\theta(1 - Q_i)$ for winner, $\varphi(1 - Q_i)$for losers. (Any firm with the highest possible cost would choose $Q_i = 1$, hence making expected surplus of zero as required for the revenue equivalence theorem to hold.) So this auction yields $(\theta - \varphi)C/4$ by direct analogy with (iib). In addition, all three firms also make $\varphi(1 - c_i)$ automatically. So expected industry profit $= 3\varphi(1 - C/2) + (\theta - \varphi)C/4$.

(iv) losing the ability to price discriminate means firms no longer rip-off their locked in consumers so much - but it also dulls competition for the uncommitted consumers. In this example the effects cancel. More generally, either could dominate.

[See Klemperer (*Review of Economic Studies*, 1995, pp. 515-539).]

The model of the question is from appendix B of Bulow and Klemperer (*Brookings Papers*: *Microeconomics,* 1998, pp. 323-394) which applies this model to the cigarette market.

A12. (*Note that those in charge of the exam slightly altered the wording of the question.*)

i. a. A is indifferent about winning at price p if $\alpha z_A + (p/k_B)(1-\alpha) = p$, or $p = k_B \alpha z_A/(k_B + \alpha - 1)$, so A stays in the bidding until the price reaches $p = k_A z_A$, in which $k_A = k_B\alpha/(k_B + \alpha - 1)$, and then quits.

b. When $\alpha = 1$, we have pure private values so A bids up to its private value. When $\alpha < 1$, we have common-value elements to valuations, so A faces a winner's curse and therefore bids less, the more aggressively that B bids (i.e., the larger is k_B). [One way of thinking about this equilibrium is that it illustrates the "strategic substitutes" and "strategic complements" analysis developed in Bulow, Geanakoplos and Klemperer (*JPE*, 1985). With pure private values, your bidding is unaffected by how aggressive your opponent is. But with a common values element, the more aggressively your opponent bids, the worse is your winner's curse, so the less aggressively you will want to bid. That is, we have strategic substitutes. And this effect is greater, the greater is the common value element, that is the smaller is α.]

c. Clearly, $k_B = k_A\alpha/(k_A + \alpha - 1)$. It is easy to see that $k_A = k_B = 1$ is a Nash equilibrium.

[there are imperfect equilibria in general (one bids to ∞, the other bids 0); assuming the z's have positive density everywhere, there are other perfect Bayesian equilibria only in the pure common values case, $\alpha = 1/2$.]

ii. a. A is indifferent about winning at price p if $\phi\{\alpha z_A + (p/k_B)(1-\alpha)\} = p$, or $p = \phi k_B \alpha z_A/(k_B + \phi\alpha - \phi)$, so A stays in the bidding until the price reaches $p = k_A z_A$, in which $k_A = \phi k_B \alpha/(k_B + \phi\alpha - \phi)$, and then quits. [If $\phi > \alpha/(1-\alpha)$, A bids ∞]

b. As before, $k_B = k_A\alpha/(k_A + \alpha - 1)$. Substituting the claimed equilibrium values of k_A and k_B into this expression for k_B, and into the expression for k_A that was derived in the previous part, confirms the claimed equilibrium.

c. $k_A = \infty$; $k_B = \alpha$; i.e., A stays in forever; B bids up only to the minimum possible value it can have conditional on its own signal. The reason is that it is common knowledge that A's value exceeds B's, so it can never be rational for B to buy if A is unwilling to.

iii a. As $\phi \to \alpha/(1-\alpha)$, B quits at αz_B so profit $= \alpha E(z_B)$. At $\phi = 1$, profit $= E\min(z_A,\ z_B)$, which is greater if z_i and z_j are independent draws from the same uniform distribution, and $\alpha < 2/3$.

b. Run a first price auction; use a reserve price; give a toehold or options or other subsidy to the weaker player (e.g., counting B's bid as ϕ times what it actually is–equivalently giving B a discount equal to a fraction $(\phi - 1)/\phi$ of its bid–yields symmetric behaviour, so profit $= ((\phi + 1)/2)\ E\min(z_A,\ z_B)$, since the auctioneer has to pay the subsidy half the time.)

A13.(*sketch answer*)

A student might begin by discussing possible designers' objectives, before addressing the two most obvious

1. efficiency:

In theory, at least, any standard auction is fine if bidders are symmetric. [This assumes there is only a single indivisible unit on offer, or bidders just want a single indivisible unit—my answer will stick to this case, since that is all that those answering the question had been taught about.]

More generally, an ascending auction does well, but needs to be "open ascending" (i.e., not just a 2nd price) if there are common value components. (With pure private values, a 2nd price auction is efficient.)

2. revenue

Begin by invoking the Revenue Equivalence Theorem (section 2.1 of 2011 revision of class notes) under the conditions of which auction design is irrelevant. (One approach would be to spend some time of explaining this result.)

Note that the independent signals assumption is less natural for common values and for private values. It's not clear how much this matters. (I hadn't discussed affiliation with those answering the question, but in any case it is not clear that its affect is quantitatively large either in practice or in theory.)

With either private or common values, an ascending auction faces the "almost-certain winners" problem. I emphasise this (and it is perhaps particularly important) for the "almost common value case" (section 3.5 of 2011 revision of class notes). It should be explained why the RET fails in this case—the reason is that a bidder's probability of winning (as a function of its "type") differs greatly across auctions.

A14. (i) Each object is sold by a separate ascending auction. (Objects may differ from each other.)

All auctions take place simultaneously.

No auction closes until no one wants to bid again in any auction.

A bidder who is the current high bidder on an object cannot remove his bid unless his bid is topped by another bidder.

There may be additional rules about the number of objects you can be the high bidder on at any one time; and/or about the amount by which any new bid must top any standing high bid; and/or about either excluding a bidder, or reducing the number of objects that it is eligible to win, if it has not been sufficiently active in the bidding.

(ii) An SAA is efficient when bidders (a) want at most a single unit each, and have (b) private values, and (c) no budget constraints. Because assuming (a) and (b), "straightforward" bidding is a Nash equilibrium. ("Straightforward bidding" means bidding in every round as if the current round is the last.) And assuming (b), (c) and "straightforward" bidding, we are guaranteed the socially efficient outcome.

(a) is necessary, because otherwise a bidder may gain from demand reduction to end the auction at a lower price than if he continued to bid "straightforwardly" .

(b) is necessary, because otherwise a bidder may gain from distorting its bidding to hide information from rivals

(c) is necessary, because otherwise a bidder may not be able to outbid a lower-value but less-budget-constrained rival (just as in an ordinary auction for a single object).

[A candidate could comment , perhaps in part (iii), that straightforward bidding is easily understood to be an optimal strategy (assuming (a), (b), and others do likewise) so bidders' behaviour is likely to approximate this, and hence to achieve efficiency in practice (not merely in theory).]

(iii) If there are either common-value components to valuations, or complementarities, there may be inefficiencies. However, even though the auction may not work well, it may nevertheless work better than alternatives in these circumstances – because the multi-round aspect allows bidders to infer common value information from others, and also to work out what complementary bundles they may be able to win.

Perhaps most important is that multiunit demand can create a variety of problems. First, it permits demand reduction ("oligopsony behaviour"), and facilitates collusion and/or predation through signalling behaviour. Second, if there are complements, bidders may face "exposure" problems.

Problems may be created by the time taken by the auction, and (relatedly) the costs of running and of participating in the auction, which might also reduce entry into the auction. Weaker bidders may be deterred from entering, as by any ascending auction design.

[A candidate could comment on the implications of revenue being an objective. There is a general presumption that efficient allocation across bidders is likely to increase revenue by creating a larger pie to share out (more precisely, greater efficiency tends to increase the likelihood of selling to bidders with higher marginal revenues). An (inefficient) reserve price might also increase the SAA's

revenues. One could make some standard comments about the roles of risk aversion, affiliation etc., in ascending vs. sealed-bid style auctions. With multiunit demand, a seller interested in revenue maximisation may do better by bundling.]

[Given the way I taught the course, candidates may choose to talk about circumstances under which core-selecting and/or product-mix auctions might perform better. But I (deliberately) wrote the question in a way that did not ask for this.]

A15. (*sketchy answer*)

With private values and no budget constraints, it is a dominant strategy to "tell the truth" and the outcome is then fully efficient!

But there ARE incentives for "dishonest" bidding if there are budget constraints, common values, or "wider-game" issues (you want to keep your information secret for reasons beyond the current auction (Rothkopf et al, JPE, 1990).

Collusion may be much easier since with complements it requires only two participants to profitably collude (by contrast with most industrial organisation context where collusion only is very effective if all market participants engage in it). Relatedly, there are unusual incentives not only to merge, but also to demerge (with complements), and shill-bid (Ausubel-Milgrom, in Cramton et al (ed), 2006).

Revenue may be lower than in other auctions.

Prices can seem odd, and "unfair" (e.g., more-aggressive/higher-value bidder pays less). Relatedly, and crucially, the auction is hard for participants to understand.

Levin and Skrzypacz's, and Janssen and Karamychev's, current working papers develop the point that since your own bids only affect other people's payments (not your own), there is a problem if bidders care about others' payments. It is a worse problem if the Vickrey auction is the final stage of multistage process (such as the recent CCAs) so the effects feed back into early-stage incentives.

For a single unit, in which case it's a second price auction, and very like an ascending auction, it may do fine–especially if the bids can be kept secret. For multiple homogeneous substitute objects, it may also be fine–and could if desired be implemented as in Ausubel (AER 2004). More generally, for auctions with substitutes it is perhaps okay (at least when the Levin-Skrzypacz-Janssen- Karamychev critique is not too problematic).

For an auction with complements there are many more problems–a sealed-bid core selecting package auction may be an improvement (Day and Milgrom, Intl Jnl of Game Theory, 2008; Erdil-Klemperer, Jnl of European Econ Assoc, 2010; etc.), perhaps after a clock auction as a prior stage (see, e.g., Maldoom, Working Paper, 2007; Cramton, Rev. Ind. Org., 2013)—though this may accentuate the Levin-Skrzypacz-Janssen- Karamychev problems.

A16. (i) v_1

(ii) prob of winning = $[((v_1 - 5)/10)^2]$; expected gain conditional on winning= $[(v_1 - 5)/3]$

Expected gain $=[((v_1 - 5)/10)^2][(v_1 - 5)/3] = [(v_1 - 5)^3]/300$

(iii) risk neutrality (this suffices for us to be able to apply the Revenue Equivalence Theorem)

[might some students talk about rational behaviour, etc? This would be reasonable as to (iii) why differs from (i)—even more so if they also worried about this in part (ii)–so maybe partial credit]

(iv) $v_1 - [(v_1 - 5)/3] = 5 + (2/3)(v_1 - 5)$ (by RET, gain is the same conditional on winning)

(v) risk averse bidders all bid higher (to take a larger probability of a smaller gain) and so gain less; risk seeking bidders all bid less aggressively and gain more

(this was not discussed in class, but was done in the problem set).

A17.

(a) A bidding strategy for player i is a function $g_i : [0,1] \to \mathbb{R}^+$ that assigns a bid for each valuation $v_i \in [0,1]$.

(b) The expected utility for bidder i with value v_i is:

$$E(u_i(g_i, g_{-i}| v_i)) = \Pr(g_{-i}(v_{-i}) < g_i(v_i)) \cdot (v_i - g_i(v_i))^{\alpha_i}$$

(assuming that draws have zero probability.)

(c) Suppose bidder $j \neq i$ uses a linear bidding strategy (bids according to $g_j(v_j) = cv_j$). The expected utility of bidder i given value v_i and bid x is:

$$E(u_i(g_i, g_{-i}| v_i, x)) = \Pr(g_{-i}(v_{-i}) < x) \cdot (v_i - x)^{\alpha_i} = \frac{x}{c} \cdot (v_i - x)^{\alpha_i} \text{ if } x \leq c.$$

FOC for the optimal x (which maximizes the expected payoff):

$$0 = \frac{\partial}{\partial x}\left(\frac{x}{c} \cdot (v_i - x)^{\alpha_i}\right) = \frac{1}{c} \cdot \left((v_i - x)^{\alpha_i} - \frac{x \cdot \alpha_i}{(v_i - x)^{(1-\alpha_i)}}\right) \Rightarrow x = \frac{v_i}{1+\alpha_i}$$

If $x > c$, bidder i will win the object for sure and should therefore never bid anything larger than c. So the equilibrium strategy is:

$$g_i(v_i) = \min\{\frac{v_i}{1+\alpha_i}, c\}$$

(d) Part C implies that there is a Bayesian-Nash equilibrium in which for each player i: $b_i(v_i) = \dfrac{v_i}{1+\alpha_i}$ if $\alpha_i = \alpha_j$, because each player is making a best reply to his opponent's strategy in this case.

(e) W.l.o.g. let $\alpha_1 < \alpha_2$. Then $b_i(v_i) = \dfrac{v_i}{1+\alpha_i}$ is not an equilibrium, because Part C implies that player 1 would do strictly better by reducing his bid to $\min\{\frac{v_1}{1+\alpha_1}, \frac{1}{1+\alpha_2}\}$. (The question did not ask this, but $b_i(v_i) = \min\{\frac{v_i}{1+\alpha_i}, \frac{1}{1+\alpha_j}\}$ is also not an equilibrium if $\alpha_i \neq \alpha_j$: if, w.l.o.g., $\alpha_1 < \alpha_2$, player 2 bidding $\frac{v_2}{1+\alpha_2}$is a best reply to $b_1(v_1) = \frac{v_1}{1+\alpha_1}$, but is *not* a best reply to $b_1(v_1) = \min\{\frac{v_1}{1+\alpha_1}, \frac{1}{1+\alpha_2}\}$. The reason is that player 1 would then be bidding exactly $\frac{1}{1+\alpha_2}$ with strictly positive probability, and player 2 would then do better to move any bids that would have been just below $\frac{1}{1+\alpha_2}$ to just above $\frac{1}{1+\alpha_2}$.)

(f) In order to justify the solution of part (d), it suffices that the distribution of values is common knowledge among the players, and, in addition, each player anticipates that the other player (j) is choosing a linear strategy $g_j(v_j) = cv_j$ for some $c \geq \frac{1}{1+\alpha_i}$(not necessarily the equilibrium strategy of part (d)).

(*I believe this is the only time that an Oxford MPhil exam question on auctions has not been set by me. The original version was flawed, and the substitute version on this website is included so that students with access to the original understand the error.*)

A18. *[sketch of possible answer]*

In both cases bidders bid demand schedules. [The auctioneer computes the aggregate demand, and finds the price at which its supply equals aggregate demand. All demand above that price is satisfied. Winning bids are filled at that price in the uniform-price auction, but at the prices actually bid in a discriminatory auction (which is therefore also sometimes called a "pay-your-bid" auction).]

A uniform-price auction bidder's demand schedule starts at the bidder's actual valuation for his first unit (exactly as in a second-price auction, and for the same reason), but a bidder for any larger quantity than a single unit recognises that buying an additional unit drives up the market price, and so drives up the price it pays for all the other units it buys. It therefore cuts back the price it bids for larger quantities exactly as a monopsonist does. (Note, however, that a bidder for a sufficiently small quantity recognises that it will not significantly affect the market price, so can simply bid its actual value for each unit it demands - as in a second-price auction.)

By contrast, a bidder's discriminatory demand schedule starts below (typically well below) the bidder's actual valuation for his first unit. Exactly as in a first-price auction, bids for high-valued units will be shaded down a long way towards the expected market clearing price. However, also as in a first-price auction, bids for low-valued units will hardly be shaded at all.

Indeed the pressure towards flat demand schedules is even greater for (multi-unit) discriminatory auctions than in (single-unit) first-price auctions. A bidder recognises that winning a second unit requires beating more opponents' bids than winning just one unit, so the bidder faces stronger competition for the second unit than for the first. (Indeed, if the bidder's "true demand" is fairly flat (i.e., it has a fairly constant marginal valuation of units), it might even want to bid more for the second unit if the rules of the auction permitted that. However, we assume the auctioneer takes bids in descending order.)

So discriminatory demands will generally be flatter than uniform-price demands, and uniform-auction bids are more dispersed.

[A good answer could provide more detail. I handed out, and discussed in the lecture, a series of pictures illustrating these points, and those below.]

(i) At least in theory, a uniform-price auction can sustain implicitly collusive strategies as a one-shot Nash equilibrium. This may particularly be a problem in frequently-repeated settings, with small numbers of bidders, difficult entry, and very stable demand (e.g. electricity). (This is probably not a problem if supply can be varied by the auctioneer, or even just if supply is elastic.) The discriminatory auction is in principle unproblematic, though if it does discourage entry it may facilitate collusion simply by reducing the number of participants.

(ii) A uniform-price auction is attractive to small bidders, because of the way in which large bidders reduce their demand. Furthermore a small bidder is protected by the fact that everyone pays the same price. (A possible exception is that weak bidders may find entering a uniform-price auction unattractive, just as weak bidders find ascending auctions of small numbers of units in which bidders can win at most one unit each unattractive.)

By contrast, in a discriminatory auction, bidders are playing "Guess the price" – that is, information about the likely market clearing price is very valuable, and collecting and analysing information is disproportionately costly for small bidders. So the uniform-price auction will increase participation. [This is especially true with common values, but the question specified private values in order to reduce the number of issues to think about.]

A good answer could note that this problem can be resolved by allowing non-competitive bidding in a discriminatory auction. Furthermore participation may not be important if there is a good resale market (provided collusion can be avoided).

(iii) Efficiency will be impacted in obvious ways by participation and collusion (discussed above).

So let's assume similar participation and no collusion: either kind of auction can then be more efficient than the other, because different kinds of bidders shade their bids downwards by different amounts (see first part), and these amounts are also different for different units.

If all bidders have flat true demands (constant marginal valuation of units) for the same number of units, the bids in a discriminatory auction are flat. (Each bidder then faces the same competition for the first unit it would win, or the second unit, or the third unit, etc. ... so bids the same for each unit.) So if their valuations are drawn from the same distribution, then just like in an ordinary single-unit auction the person with the highest valuation bids the most, in Nash equilibrium, so we have full efficiency! By contrast, a uniform price auction is not fully efficient since each bidder bids a set of decreasing prices for units it values the same.

However, if each bidder has downward sloping demand, the discriminatory auction may be less efficient than the uniform price auction even if everybody's valuations are drawn from the same distribution. (In a discriminatory auction for m units, a bidder's bid to win an nth unit depends upon its nth-highest valuation, which is generally drawn from a different distribution than the distribution of the $m-n+1$th opponent's valuation (which determines the bid it is competing against).)

(Ausubel, Cramton, Pycia, Rostek, and Weretka "Demand Reduction and Inefficiency in Multi-Unit Auctions", REStud, 2014 [http://pycia.bol.ucla.edu/ausubel-cramton-pycia-rostek-weretka-auctions.pdf], provide an example (Ex IV, p1376) in which the uniform price auction is more efficient whenever bidders' values of winning a second unit are a sufficiently small fraction of their values for winning a first unit.)

Moreover, the discriminatory auction is not efficient for bidders for a single unit if their values are drawn from different distributions.

By contrast, the uniform-price auction is perfectly efficient when bidders only demand a single unit, and generally also close to this when bidders are all "small".

More generally, if bidders are not too large, bidding is probably easier in the uniform-price auction, so actual bidding may more closely approximate the theoretical equilibrium, and outcomes may be more efficient.

A very impressive answer might observe that some other issues may affect efficiency more broadly than in the immediate outcome of this auction: if bidders' demands are not too large relative to the auction size, uniform bidding is likely to be more informative; on the other hand, because information is less valuable to bidders in a uniform auction, bidders may collect more information before a discriminatory auction, so bidding and allocation, especially in the longer run, may be more efficient (ignoring resources "wasted" collecting privately-useful but socially-useless information, and the fact bidders may strategically withhold information so any pre-market will be thinner and less informative ...). (There may be other issues too: a uniform price auction makes a short squeeze harder; possibly effects of budget constraints, risk-aversion, asymmetries, etc.)

A19. [Sketch of a possible answer; a good answer was one that was relevant to the context given, rather than just discussing Vickrey auctions in general.]

(a &b) With private values (as the question assumed) and no budget constraints, etc., it is a dominant strategy to "tell the truth" and the outcome is then fully efficient.

However, there are incentives for "dishonest" bidding if there are "wider-game" issues (bidders want to keep their information secret for reasons beyond the current auction, see, e.g.,Rothkopf et al, JPE, 1990), and/or there are budget constraints.[8] In this very simple case, a bidder with a value exceeding its budget constraint will simply bid its budget constraint, so in that sense bidding is honest, but the outcome may not be socially efficient.

[8]There would also be incentives for "dishonest" bidding if there are common values, but I deliberately abstracted from the latter. I also deliberately abstracted from relative performance issues, to minimise Janssen-Karamychev/Levin-Skrzypacz-style problems that since your own bids only affect other people's payments (not your own), there is a problem if bidders care about others' payments, and a worse problem if the Vickrey auction is the final stage of multistage process (such as the recent CCAs) so the effects feed back into early-stage incentives.

Revenue may be lower than in other auctions—especially in this complements setting.

Prices can seem odd, and "unfair",[9] and the auction may be hard for participants and observers to understand, even in a setting this simple.

So the merits of a Vickrey auction depend on the auctioneer's objectives, including the extent to which it cares about participants' and observers' perceptions and behaviour.

Collusion may be much easier than with other designs or in other settings since, in this complements context, the Vickrey auction allows only two participants to profitably collude (by contrast with most contexts where collusion only is very effective if all market participants engage in it). Relatedly, there are unusual incentives not only to merge, but also to demerge [more precisely, in this case, to stay de-merged] (since we have complements), and shill-bid (see, e.g., Ausubel-Milgrom, in Cramton et al (ed), 2006). [My lecture gave examples relevant to the setting of this problem.] These problems are likely to be particularly severe if the bidders know the number and identities and interests of the other bidders, and if co-operation and/or (implicit or explicit) collusion is feasible–so these factors are crucial to the merits of a Vickrey auction, since there is such a tiny number of bidders.

(c) Since we don't have common values, and there's also no coordination issue in a problem this simple, there would be limited benefit to any first-stage dynamic process here. Indeed a first-stage may actually aggravate problems (e.g., it might facilitate collusion or even allow someone to risk pushing up price on the object it does not want to exhaust the complement-bidder's budget).

But a sealed-bid core-selecting package auction may be an improvement. See, especially, Day-Milgrom, in Neeman et al (eds.), 2013 [this paper updates and corrects their earlier Intl Jnl of Game Theory, 2007 paper]; smaller contributions include Erdil-Klemperer, Jnl of European Econ Assoc, 2010; etc.

(There are, of course, other possibilities too (e.g., a revenue-maximising auctioneer might seek to exploit risk-averse bidders), but it is worth noting that unless we choose some kind of package format, optimal, or even "sensible", bidding will be very hard for the complements bidder.)

[9] With substitutes, a more-aggressive/higher-value bidder often pays less; in our context, the issue is the difference in likely payments between the alternatives of the combination of East and West winning versus the "pair" bidder winning.

A20. (i) In the single-unit case, the auctions are strategically identical.

Even if they are strategically identical, they may be played differently in practice. (The Nash equilibrium is not dominant strategy.)

[Although we did not discuss this in the lectures, it could be mentioned that with multiple units they are not strategically identical (and nor are equilibrium outcomes identical, though they may—-under tight conditions—be revenue equivalent).]

(ii) If the ascending auction is a Japanese auction for a single unit, the auctions are strategically identical if there are just two players. [A good answer would explain the rules of a Japanese auction, and how the ability to use complicated strategies, make jump bids, etc makes a free-form ascending auction different.]

With private values (and any number of bidders), a Japanese auction OR an ascending auction (free-form, but ignoring the bidding increment) is identical in equilibrium outcome (and likely to have the same outcome, because the Nash equilibrium is in dominant strategies), but is not strategically identical, to a second-price auction.

[Since I asked about second-price auctions, I would not expect students to address multi-unit issues (and, anyway, we did not discuss them in the lectures), but it could be added that (with private values) a (k+1th) price auction for k units (with bidders who want 1 unit each) is identical in equilibrium outcome (and likely to have the same outcome), but is not strategically identical, to a Japanese auction or an ascending auction.]

A good answer might note that in an ascending auction the winning bidder will not (in general) reveal her actual value. In practice, bidders may be unwilling to bid their true values in a second-price auction for fear of revealing information that may be useful to others (such as regulators, competitors, or those they might have to bargain with such as suppliers) after the auction.

(iii) When the assumptions of the revenue equivalence theorem hold they all yield the same expected payment by each bidder and the same expected revenue for the auctioneer, in Nash equilibrium.

There are various forms of this theorem. This was the one I proved in the lecture notes:

Revenue Equivalence Theorem (IPV Case): *Assume each of n risk-neutral potential buyers of an object has a privately known value independently drawn from a common distribution $F(v)$ that is strictly increasing and atomless on $[\underline{V}, \overline{V}]$. Any auction mechanism in which (i) the object always go to the buyer with the highest value, and (ii) any bidder with value $\underline{V}$ expects zero surplus, yields the same expected revenue, and results in a buyer with value v* making the same expected payment.

[This one (also in the lecture notes) would be even better:

Assume each of n risk-neutral potential buyers of an object has a privately known signal independently drawn from a common distribution $F(z)$ that is strictly increasing and atomless on $[\underline{Z}, \overline{Z}]$. Any auction mechanism in which (i) the object always go to *the buyer with the highest signal, and (ii) any bidder with signal $\underline{Z}$ expects zero surplus, yields the same expected revenue, and results in a buyer with signal z making the same expected payment.*]

The key assumptions are *risk-neutrality*, *independent signals (or values),* and *common distribution (i.e., symmetry).* (Also *given number of bidders*, and *non-cooperative equilibrium bidding.*)

sketch of proof of IPV case

Consider *any* mechanism for allocating the object; let $S_i(v_i)$ be the expected surplus that bidder i will obtain in equilibrium; let $P_i(v_i)$ be her probability of receiving the object in the equilibrium.

Then $S_i(v) \geq S_i(\widetilde{v}) + (v - \widetilde{v})P_i(\widetilde{v})$

The right-hand side is the surplus that i would obtain if she had type v but deviated from equilibrium behaviour, and instead followed the strategy that type $\widetilde{v}$ of i is supposed to follow in the equilibrium of the game induced by the mechanism. That is, if type v exactly mimics what type $\widetilde{v}$ would do, then v makes the same payments and wins the object as often as $\widetilde{v}$ would. So v gets the same utility that $\widetilde{v}$ would get $(S_i(\widetilde{v}))$, except that in states in which $\widetilde{v}$ would win the object (which happens with probability $P_i(\widetilde{v})$) type v values the object at $(v - \widetilde{v})$ more than $\widetilde{v}$ does, and so v obtains an extra $(v - \widetilde{v})\ P_i(\widetilde{v})$ more surplus in all.

So, since type v mustn't want to mimic type $v + dv$,

we have $S_i(v) \geq S_i(v + dv) + (-dv)P_i(v + dv)$.

And since $v + dv$ mustn't want to mimic type v,

we have $S_i(v + dv) \geq S_i(v) + (dv)P_i(v)$.

Reorganising, $P_i(v + dv) \geq \frac{S_i(v+dv)-S_i(v)}{dv} \geq P_i(v)$,

and taking the limit as $dv \to 0$ we obtain $\frac{dS_i}{dv} = P_i(v)$.

Integrating up,$S_i(v) = S_i(\underline{V}) + \int_{x=\underline{V}}^{v} P_i(x)dx$.

So any mechanisms which have the same $S_i(\underline{V})$ and the same $P_i(v)$ for all v and for every i, have the same $S_i(v)$. So any given type of player i makes the same expected payment in each of the mechanisms (since $S_i(v) = vP_i(v) - E$(payment by type v of player i), since the bidder is risk-neutral). So i's expected payment averaged across her different possible types, v, is the same for the mechanisms. Since this is true for all bidders, the mechanisms yield the same expected revenue for the auctioneer.

In particular any mechanism (including all four of our auctions) which always gives the object to the highest-value bidder in equilibrium (so $P_i(v) =$ the probability that all $(n-1)$ other bidders have lower values than i), and give a bidder of the lowest feasible type no surplus (so $S_i(\underline{V}) = 0$) yield the same expected payment by each bidder and the same expected revenue for the auctioneer.

A.21 (i) A,B,C bid 5,7,10 for the packages they want, so A,B win and pay 3,5

(ii) (a) Vickrey-nearest: same winners, but they pay 4,6 (since total must be 10);

(b) Reference rule (Erdil and Klemperer, *JEEA*, 2010): $3\frac{1}{3}$, $6\frac{2}{3}$, since objects are identical and B wins two (the objects are identical, so the prices should be as near as possible to identical, but with total=10.

[other rules could be used, if justified by the examinee]

(iii) A,B',B",C bid 5,3,4,10 so A,B',B" win. They pay 3,1,2 in Vickrey;

they pay $4\frac{1}{3}$, $2\frac{1}{3}$, $3\frac{1}{3}$ in Vickrey-nearest; and they pay $3\frac{1}{2}$, 3, $3\frac{1}{2}$ in Reference rule (since we can't charge more than 3 to B', so have to charge $3\frac{1}{2}$ to A,C)

(iv) it may be natural to (a) compare the Vickrey rule against non-Vickrey core-selecting-package alternatives in general; then (b) discuss distinctions between the non-Vickrey alternatives:

(a) The alternatives are intended to reduce the Vickrey auction's problems that revenue may be very low (as in the example—since a loser was prepared to pay 10), collusion is easier than with other designs in complements contexts such as this one (since it allows only two participants to profitably collude, by contrast with most contexts where collusion only is very effective if all market participants engage in it), and the related unusual incentives to merge, or demerge (as in the example) (or stay de-merged), and/or shill-bid (see, e.g., Ausubel-Milgrom, in Cramton et al (ed), 2006). [My lecture gave examples relevant to the setting of this problem.] (The question did not ask for a general discussion of the merits of a Vickrey auction.)

However, it is very hard to understand how to bid sensibly in any core-selecting rule even in many simple settings. And if bidders can figure out how to bid, Jacob Goeree and Yuanchuan Lien (*Theoretical Economics* 2016, On the Impossibility of Core-Selecting Auctions) shows that with equilibrium bidding, core-selecting auctions may well result in lower than Vickrey revenues, and in inefficient outcomes that are on average further from the core than Vickrey outcomes! (They argue that no Bayesian incentive-compatible core-selecting auction exists.)

(b) Reference rules seem fairer and more comprehensible and Erdil-Klemperer showed they are likely to give lower marginal incentives to bidders to deviate from truth-telling (and probably suffer less from "odd" incentives to merge, demerge, collude and shill than Vickrey-nearest (as well as than Vickrey)). However, the choice of reference prices may be contentious if objects are not identical. Erdil-Klemperer give a simple example in which Vickrey-nearest is ex-ante welfare maximising among all minimum revenue core-selecting rules, but Larry Ausubel and Oleg Baranov's (*mimeo* 2013, Core-Selecting Auctions with Incomplete Information) results suggest the opposite. More generally, Ausubel-Baranov looks at four different core-selecting auction formats suggested in the literature, and seem to like the proxy rule of Ausubel-Milgrom which is a form of Erdil-Klemperer reference rule, for revenues and efficiency, especially as correlations among bidders' values increase.

A.22 *[N.B.: I do not expect students to know (or write) anything like as much detail as I have given here.]*

0. This is a "short" question, so I don't think a good answer requires descriptions of the different auctions. But a student with less to say might tell us what they are: Both SMRAs and PMAs transact multiple potentially-different objects. For an auction that is selling goods, a SMRA is a simultaneous ascending auction, and is described in the answer to my 2012 question (auctions to buy are the opposite). PMAs are described at www.nuff.ox.ac.uk/users/klemperer/productmix.pdf.

1. Given the way the question is posed it would not be necessary, but it might be natural, to explain that when bidders (a) want at most a single unit each, (b) have private values, and (c) have no budget constraints, "straightforward" bidding is a Nash equilibrium of both the SMRA and the PMA. (Strictly speaking we also need that the bidding increment is infinitesimal in the case of the SMRA.) That is, in the SMRA each bidder bids in every round as if the current round is the last; in the PMA each bidder simply expresses his true values.

Moreover, "straightforward" bidding is also optimal in both the SMRA and the PMA for any bidder who wants multiple objects but thinks his demand is too small to affect market prices, and valuations are strong substitutes (i.e., all units of all goods are mutual substitutes), for all bidders.

So in these cases the SMRA and the PMA are equivalent (and efficient).

2. The distinctions between the auctions flow from the fact that the SMRA is dynamic (multi-round) and the PMA is static (single-round):

(a) *advantages of the SMRA:*

(i) If there are either common-value components to valuations, or complementarities (including any non-strong-substitutabilities), neither the SMRA nor a simple PMA allows bidders to accurately express their preferences in general, and neither is efficient in general. (PMAs in which bidders bid ordinary "dot bids" restrict bidders to expressing strong-substitute preferences, so force bidders to approximate any other preferences by strong-substitute preferences.)

The SMRA may work better than the PMA — because the SMRA's multiple rounds allows bidders to infer common-value information from others. (However, the PMA could be extended to allow bidders with "common values" to revise their bids based on reported "interim" auction prices.)

The SMRA's multiple rounds can also allow bidders to learn something about which objects they are likely to win and then focus their attention on the values of those, without having to waste resources analysing their valuations for objects that they would have no chance of winning.

For the same reason, the SMRA helps "package discovery", that is, helps bidders work out what complementary bundles they may be able to win, and so might sensibly bid for.

(ii) Another advantage of the SMRA is that winning bidders need not (in general) reveal their actual values. So the SMRA has the same advantage over the PMA that the single-unit ascending auction has over a single-unit second-price auction–bidders may be unwilling to bid their true values in PMA for fear of revealing information that may be useful to others, such as regulators, competitors, or suppliers they might have to bargain with after the auction. (However, a PMA can be run in a way that no-one sees the winning bids.)

(iii) Bidders with budget constraints will generally find bidding easier (and so outcomes maybe more efficient) in an SMRA than in a standard PMA (though there is a version of the PMA for budget constrained bidders, which may work better than either an SMRA or a standard PMA in this context).

(b) *advantages of the PMA:*

(i) A potentially major advantage of the PMA is that the multi-round aspect of the SMRA may allow bidders to coordinate their bidding, possibly facilitated by the ability to signal. Bidders in an SMRA may be able to implement collusion (as illustrated, for example in the lectures' examples of the Austrian and German auctions), or some bidders may be able to follow predatory strategies (see, for example, the lectures' example of an early US auction).

Recent versions of the SMRA are more robust than earlier versions, but SMRAs remain vulnerable to this kind of behaviour. (In principle, implicit collusion might be a Nash equilibrium of a PMA, just as it is in an ordinary uniform price auction, but this seems even less likely than in most uniform-price-auction settings, because of the extreme difficulty of coordinating behaviour across differentiated goods.)

[Less important, for bidders who want more than one unit, it is generally individually rational (i.e., ignoring coordination possibilities) to engage in demand reduction. In experiments in auctions of a single variety, they do this much less than equilibrium analysis would imply (at least if there are more than two bidders and/or more than two units available) but they also do this much more in a dynamic auction than in a static auction. So non-coordinated demand reduction may be less in a PMA, which would make it more efficient (although non-coordinated demand reduction is not always a significant problem in an SMRA either). For various reasons competitive equilibrium seems a good assumption in the Bank of England's PMA, and likely more generally.]

(ii) The fact that the PMA's outcome is found by solving a static problem based on stated preferences has several advantages relative to the SMRA's approach. The latter relies on bidders revealing the necessary information over time in response to monotonically increasing prices and also requires activity rules. (Absent activity rules, bidders in an SMRA typically have incentive to play ”snake in the grass” strategies of waiting for others to reveal their hands first; even absent common values complementarities, and budget constraints, a bidder may gain from waiting for reason (aii) above.)

First, when rationing is permitted, versions of the PMA can, unlike the SMRA, in theory find the efficient solution when participants have any (ordinary) substitutes valuations (these are represented using “asymmetric dot bids”) and/or some kinds of complements valuations. (The Bank of England's 2014 implementation includes complementarities on the seller's side.)

Second, the PMA allows the auctioneer to easily express its own preferences about how the mix, and the total quantity, of objects sold should depend on the bidding, and the PMA also extends naturally to permitting multiple sellers as well as buyers. (The SMRA could be modified to do these things, but I am not aware of this having been suggested prior to Klemperer (2008). Moreover, the PMA's approach to incorporating auctioneer preferences is simpler and more intuitive than a modified SMRA would be.)

Third, a PMA can easily implement objectives different than efficiency, e.g., auctioneer revenue maximisation, and it is possible to make other modifications such as varying the pricing rule. (These things don't seem possible in anything that would be recognisable as a version of an SMRA.)

(iii) The speed of the PMA makes it less costly for participants. It may therefore also increase participation. (The SMRAs multiple-round bidding may also be impractical in, e.g., financial markets with rapidly changing conditions.)

A.23 *[sketch of possible answer]*

0. In a discriminatory auction winning bids are filled at the prices actually bid– exactly as in a first-price auction.

In a uniform price auction, winning bids are filled at a price where demand equals supply. In economists' models, this is usually the highest losing bid–exactly as in a second-price auction. (In reality the price chosen is usually the lowest winning bid, but with many units, this is usually approximately, or exactly, the same price.)

1. *Assuming a fixed number of bidders who play Nash equilibrium strategies*,

a uniform-price auction bidder's demand schedule starts at the bidder's actual valuation for her first unit (exactly as in a second-price auction, and for the same reason), but a bidder for any larger quantity than a single unit recognises that buying an additional unit drives up the market price, and so drives up the price it pays for all the other units it buys. It therefore cuts back the price it bids for larger quantities exactly as an oligopsonist does.

By contrast, a bidder's discriminatory demand schedule starts below the bidder's actual valuation for her first unit. As in a first-price auction, bids for high-valued units will be shaded down towards the expected market clearing price. However, also as in a first-price auction, bids for low-valued units will hardly be shaded at all. Indeed the pressure towards flat demand schedules is even greater for (multi-unit) discriminatory auctions than in (single-unit) first-price auctions. A bidder recognises that winning a second unit requires beating more opponents' bids than winning just one unit, so the bidder faces stronger competition for the second unit than for the first. (Indeed, if the bidder's "true demand" is fairly flat (i.e., it has a fairly constant marginal valuation of units), it might even want to bid more for the second unit if the rules of the auction permitted that. However, we assume the auctioneer takes bids in descending order.)

So:

1a. If bidders want only a single unit each, the analogy is a very good one:

Bidders "tell the truth" in the uniform case, just as in a second-price auction, but shade below their values in the discriminatory case, just as in a first price auction.

And uniform price auctions will generally be more efficient than discriminatory auctions for the same reason (neither the discriminatory auction nor the first-price auction is efficient for bidders for a single unit if their values are drawn from different distributions).

Discriminatory demands will generally be flatter than uniform-price demands, and uniform-auction bids are more dispersed (just as second-price bids will be more dispersed than first priced bids in a single unit auction).

Moreover, in this special case, revenue equivalence holds in the multiunit case as in the single unit case.

1b. The analogy remains pretty good for bidders whose demands are "small" relative to the total demand. (A bidder for a sufficiently small quantity in a uniform-price auction recognises that it is very unlikely to significantly affect the price [so it can simply bid its actual value for each unit it demands, as in a second-price auction]; and a bidder for a sufficiently small quantity in a discriminatory expects to face a similar strength of competition across the range of quantities it might win.)

1c. When bidders are not "small", analogies with the single unit auction cases become less helpful, as the discussion above makes clear.

Either kind of auction can then be more efficient than the other, because different kinds of bidders shade their bids downwards by different amounts, and these amounts are also different for different units.[10]

[10] If all bidders have flat true demands (constant marginal valuation of units) for the same number of units, the bids in a discriminatory auction are flat. (Each bidder then faces the same competition for the first unit it would win, or the second unit, or the third unit, etc. ... so bids the same for each unit.) So if their valuations are drawn from the same distribution, then just like in an

Moreover, at least in theory, a uniform-price auction can sustain implicitly collusive strategies as non-cooperative equilibria (see below). The discriminatory auction is in principle unproblematic, in this respect.

2. *If we worry about participation, and coordination/collusion possibilities*,

that is, we drop the assumption of a fixed number of bidders who play Nash equilibrium strategies, the analogies become much less good:

2a. *Participation*: A uniform-price auction is attractive to small bidders, because of the way in which large bidders reduce their demand[11]. Furthermore a small bidder is protected by the fact that everyone pays the same price. By contrast, in a discriminatory auction, bidders are playing "Guess the price" – that is, information about the likely market clearing price is very valuable, and collecting and analysing information is disproportionately costly for small bidders. So the uniform-price auction will increase participation (especially with common values). (This "Guess the price" problem can be resolved in a discriminatory auction by allowing non-competitive bidding, but that is rather like introducing a uniform price component into the discriminatory auction.)

By contrast, a first-price auction for a single unit creates more uncertainty about who might win, and so is more likely to attract entrants, than a second-price auction (especially in some common-value settings).

2b. *Coordination/collusion possibilities*: Since, at least in theory, a uniform-price auction can sustain implicitly collusive strategies as non-cooperative equilibria, it may not be too hard for bidders to sustain co-operation in more frequently-repeated settings, with small numbers of bidders, difficult entry, and very stable demand (e.g. electricity). (This is probably not a problem if supply can be varied by the auctioneer, or even just if supply is elastic.) The discriminatory auction is less problematic in this respect (though if it does discourage entry it may facilitate collusion simply by reducing the number of participants). So here there is a closer analogy with the single-unit auction formats, where it is usually argued that a second-price auction makes collusion somewhat easier.

---

ordinary single-unit auction the person with the highest valuation bids the most, in Nash equilibrium, so we have full efficiency! By contrast, a uniform price auction is not fully efficient since each bidder bids a set of decreasing prices for units it values the same.

However, if each bidder has downward sloping demand, the discriminatory auction may be less efficient than the uniform price auction even if everybody's valuations are drawn from the same distribution. (In a discriminatory auction for m units, a bidder's bid to win an nth unit depends upon its nth-highest valuation, which is generally drawn from a different distribution than the distribution of the $m-n+1$th opponent's valuation (which determines the bid it is competing against).)

(Ausubel, Cramton, Pycia, Rostek, and Weretka "Demand Reduction and Inefficiency in Multi-Unit Auctions", REStud, 2014 [http://pycia.bol.ucla.edu/ausubel-cramton-pycia-rostek-weretka-auctions.pdf], provide an example (Ex IV, p1376) in which the uniform price auction is more efficient whenever bidders' values of winning a second unit are a sufficiently small fraction of their values for winning a first unit.)

[11] See also Baisa and Burkett (2016 working paper)'s model of a large bidder with multi-unit demand vs continuum of single-unit demanders.