Efficient cost sharing with a cheap residual claimant*

August 2007

1 The residual claimant approach

The residual claimant is a simple way to align *monetary* incentives and efficiency. Its most familiar application is to the rich class of Vicrey-Clarke-Groves (VCG) mechanisms. If individual preferences are quasi-linear in money (or any other numeraire) a VCG mechanism takes the "efficient" decision, except for the fact that the budget imbalance must be covered by a residual claimant ([**?**, 7]). Thus we achieve full efficiency (Pareto optimality) in the augmented economy where the claimant's preferences are taken into account, and he only cares about cash transfers. In this case we speak of *residual efficiency* in the initial economy.

I submit that the realism of the residual claimant idea hinges around the *size* of the cash transfer he or she receives. If the residual claimant (thereafter RC) pockets a surplus commensurate to -or larger than- that available in the economy, the cost of generating the efficient decision is prohibitive: the RC receives a significant rent, and the choice of the entity playing the role of RC is a matter of dispute. These difficulties are amplified if the RC must finance a substantial deficit, in effect paying out an additional rent to the participants: now it may not even be feasible to find an entity willing to play the RC role.

I propose a canonical mechanism, called *residual**, to run a one commodity convex technology, and argue that its residual claimant is cheap for many specifications of the technology.

The residual* mechanism is simpler than a VCG mechanism: individual messages are one-dimensional "demands", namely a request for a certain amount of output, which the mechanism must meet. Its incentives properties are weaker: individual demands are typically not dominant strategies. However for any profile of convex quasi-linear preferences, the non cooperative analysis of the demand game is quite straightforward, because it is a *potential game* ([15]). Even when information on preferences is entirely private, most familiar learning algorithms such as fictitious play and best reply dynamics converge to a Nash

*This is an extended abstract of my paper "Efficient and almost budget-balanced cost sharing", available at http://www.ruf.rice.edu/~econ/faculty/Moulin.htm

Dagstuhl Seminar Proceedings 07261
Fair Division
http://drops.dagstuhl.de/opus/volltexte/2007/1231

equilibrium outcome ([15, 16, 12]). Therefore in our demand games, simple patterns of decentralized behavior by the participants achieve one of the Nash outcomes, irrespective of the informational assumptions.

For many smooth cost functions, in particular for all analytic functions, the residual* mechanism generates a budget transfer to the RC that is *vanishingly small* relative to the overall surplus in the economy, when the set of potential users of the commons grows. Although it is still easier to implement an RC who burns some money than one who must finance the deficit, the benefit of aligning efficiency with incentive-compatibility far outweighs the small cost of establishing an independent residual claimant.

One of the oldest[1] normative requirements of a resource allocation mechanism is that agents should have no objection to participate. In many contexts, this results from private property rights (e.g., in any trading mechanism); in other contexts (provision of a local public good), it reflects a concern for fairness: an agent who is forced to participate subsidizes other agents, she does not get a share of the collective surplus.

For the rich class of *totally monotone* cost functions, the residual* mechanism guarantees *voluntary participation.* In particular a user with a null or small demand may reseive a cash *subsidy.* This makes good sense when marginal costs increase: refraining from using the technology benefits all active users. In our mechanism inactive users in effect have a claim on the surplus generated by the active users. Thus the manager of the mechanism must monitor *entry*, lest many newcomers show up to demand no output and receive some cash.

Relation to the literature

Familiar applications of demand mechanisms include buyers cooperatives ([20]) and capacity sharing ([13, 1]). Recent applications are to resource allocation in networks with elastic supply,such as sharing the cost of a bandwidth, and assigning service priorities([6, 8, 9, 10, 23, 27, 4]).

Our asymptotic efficiency results (Theorems 2,3,4) differ from classic findings of the same flavour in the literature on VCG mechanisms and on double auctions. The latter are all *probabilistic* statements about the *expected* convergence of the equilibrium surplus to the efficient level. In general VCG mechanisms, this can be achieved by polling a subset of participants ([7, **?**]); in double auctions, the expected surplus in any non cooperative equilibrium outcome converges to zero ([24, 14, 5]). Thus these results are predicated on the existence of common priors on individual types. By contrast, our convergence results bear on the *worst case* configuration of types, they are completely independent of any prior belief on individual types. Evaluating the worst relative surplus loss in equilibrium, over the entire domain of preference profiles, was pioneered by the recent literature on the *price of anarchy* in congestion games ([11, 22, 21, 2, 17]).

[1] It goes back to the early literature on public finance. See e.g. Wicksell's entry in [19].

2 Overview of the results

The residual* mechanism asks each potential user i of the technology to *demand* an arbitrary non negative quantity x_i of "service" (output). Given the increasing, convex, and differentiable cost function C, for each demand profile x the mechanism computes monetary (input) charges y_i. Each individual demand is served in full by the mechanism, who can only adjust the monetary charges to the various users, including those who choose not to consume. This defining feature of simple cost sharing mechanisms is the property called Consumer Sovereignty in [18][2].

Given quasi-linear utility functions $u_i(x_i, y_i) = v_i(x_i) - y_i$ for each user, the predicted outcome is some Nash equilibrium of the resulting demand game. To guarantee *residual efficiency (*Pareto optimality in the economy augmented by the residual claimant) at all equilibria and for all utility profiles*,* we must choose $y_i = C(\sum_j x_j) - h_i(x_{-i}, C)$, where the function h_i does not depend on x_i but is otherwise arbitrary (Lemma 1). By Nash's theorem if v_i is concave the demand game has at least one Nash equilibrium, and at each equilibrium $v_i'(x_i) = C'(\sum_j x_j)$ (or $v_i'(x_i) \leq C'(\sum_j x_j)$ if $x_i = 0$) for all i. Therefore the equilibrium demand $\sum_j x_j$ is optimal, i.e., it maximizes total surplus $P(x) = \sum_i v_i(x_i) - C(\sum_i x_i)$. Moreover the function P is a *potential* for the demand game, thus ensuring strong convergence properties of the classic learning algorithms ([15, 16, 12]). Finally, when n potential users share the technology, the *budget imbalance* (surplus or deficit) $\Delta(x, C) = \sum_i y_i - C(\sum_i x_i) = (n-1)C(\sum_i x_i) - \sum_i h_i(x_{-i}, C)$ is transferred to the RC.

At the equilibrium demand profile x, $r_n(x, C, v) = \frac{|\Delta(x,C)|}{\sum_i v_i(x_i) - C(\sum_i x_i)}$ measures the relative inefficiency of the equilibrium outcome. We call this ratio the *residual cost* of the mechanism at equilibrium x.

A natural choice of the transfers functions is $h_i(x_{-i}, C) = C(x_{N \smallsetminus i})$, defining the *incremental*$^+$ mechanism where each user pays the incremental cost of adding her own demand to that of other users[3]. Note that these transfers are defined by the normatively appealing property that a null demand, $x_i = 0$, results in a zero charge, $y_i = 0$. Convexity of C ensures that the RC receives a positive *surplus*, however this surplus may exhaust all but a $\frac{1}{n}$-th fraction of the efficient surplus: for all n we can choose a profile v of utilities and a corresponding Nash equilibrium x, such that $r_n(x, C, v) \geq 1 - \frac{K}{n}$, where the constant K does not depend on n.

With the notation $x_S = \sum_{i \in S} x_i$ for any subset S of agents, our *residual* mechanism* is defined as follows

$$h_i^*(x_{-i}) = (n-1)\{C(x_{N \smallsetminus i}) - \frac{1}{2} \sum_{j \in N \smallsetminus i} C(x_{N \smallsetminus i,j}) + \frac{1}{3} \sum_{j,k \in N \smallsetminus i} C(x_{N \smallsetminus i,j,k}) -$$

[2] Compare more 'bossy' mechanism eliciting full-fledged preferences and assigning to each agent an input contribution and an output share: see e.g., [20].

[3] This idea has a long history in the cost sharing literature: see [25, 3]

$$\cdots + \frac{(-1)^k}{k+1} \sum_{T \subset N \setminus i, |T|=k} C(x_{N \setminus \{i \cup T\}}) + \cdots + \frac{(-1)^{n-2}}{n-1} \sum_{j \in N \setminus i} C(x_j)\}$$

The first remarkable property of the cost shares $y_i^* = C(\sum_i x_i) - h_i^*(x_{-i})$ is that they cover exactly the budget whenever the cost function C is a polynomial of degree at most $n-1$. In this case every Nash equilibrium of the demand game is fully efficient and a residual claimant is not needed. Polynomials of degree $n-1$ or less are in fact the *only* cost functions for which we can choose the h_i to balance the budget at all demand profiles. For all other cost functions, residual efficiency implies budget imbalance at some demand profile, and our only option is to make it small.

Second property of residual*: its budget is balanced if not all potential users are active. For any cost function C, we have $\sum_i y_i^* = C(\sum_i x_i)$ if $x_i = 0$ for some i. In combination with the symmetric treatment of all individual demands, this property characterizes the residual* mechanism.

A third remarkable property is that for many cost functions, the sign of the budget transfer to the RC is independent of the utility profile: if n users share the commons, and the n-th derivative $C^{(n)}$ is everywhere non negative (resp. everywhere non positive), the residual* mechanism always generates a budget surplus (resp. a deficit).

Two of our main results, in section 6, apply to *totally monotone* cost functions C, namely indefinitely differentiable with *all* derivatives $C^{(k)}, k = 1, \cdots$ non negative on $\mathbb{R}_+$. Such functions are analytic, $C(a) = \sum_1^\infty \lambda_k a^k$, with $\lambda_k \geq 0$ for all k. Here the residual* mechanism is compelling. On the one hand it never generate a budget deficit, and its residual cost is at most $\min\{\frac{2}{\log n}, 1\}$, irrespective of C, the utility profile $v = (v_i)$, and the Nash equilibrium x. On the other hand it guarantees the two critical normative requirements *Ranking (RKG)* and *Voluntary Participation (VP).*

Ranking is a classic test of fairness (e.g.,[**?**]), requiring cost shares to be comonotonic with demands: $x_i \leq x_j \Rightarrow y_i \leq y_j$. Absent this property, some users can claim convincingly that their charge is unfair, i.e., they pay more for less service!

Voluntary Participation is an even more common[4] normative test, with a stronger incentives flavour than RKG. If users can decline to participate in the mechanism, we cannot sustain an outcome where they are worse off than by opting out, namely receiving no output at no charge. In our model, VP means that a null demand is never taxed: $x_i = 0 \Rightarrow y_i^* \leq 0$. Note that this does not preclude to *subsidize* a null demand, $x_i = 0$ and $y_i^* < 0$. Unlike the incremental one, our residual* mechanism routinely hands out cash to *inactive* users; for instance we show that with many more inactive than active users, the rent of the former group is quite large. Recall that if there is at least one inactive user, the charges will exactly cover the costs: that is because inactive users are, in effect, sharing the residual claim. Thus the right to participate in the residual* mechanism must be carefully monitored.

[4]A formal statement in the cost sharing context is in [18].

We examine the performance of our mechanism for several other classes of smooth cost functions. The first one contains the positive *generalized polynomials,* of the form $C(a) = \sum_{k=1}^{K} \lambda_k a^{p_k}$ where $p_k > 1$ and $\lambda_k > 0$. Although it is no longer possible to sign the budget imbalance, the asymptotic efficiency property of the residual* mechanism remains strong: the *worst case* residual cost (over all utility profiles and all Nash equilibria) converges to zero when n grows as $\frac{1}{n^{\inf_k\{p_k\}-1}}$. Numerical evidence strongly suggests that both Voluntary Participation and Ranking still hold for these functions, but I have been unable to prove or disprove this conjecture.

We also consider the class of *analytic* functions $C(a) = \sum_1^{\infty} \lambda_k a^k$ where the sign of λ_k is arbitrary, and to the generalized polynomials $C(a) = \sum_1^K \lambda_k a^{p_k}$ with $p_k > 1$ and $\lambda_k \in \mathbb{R}$. Here Voluntary Participation and Ranking may fail. However a weaker form of asymptotic efficiency still holds: the residual cost converges to zero as the *set* of users increases (as opposed to an increase in the *number* of users). The convergence is exponential in the former case, and hyper-polynomial in the latter.

We show finally that the regularity of the cost function is critical to the good behavior of our mechanism. For a piecewise linear cost function, and more generally for a non differentiable cost function, the residual* mechanism may generate surpluses *and* deficits growing exponentially with n. It grossly violates VP and RK. Moreover, for a piecewise linear cost *any* residually efficient mechanism (i.e., with cost shares of the form $y_i = C(\sum_i x_i) - h_i(x_{-i}, C)$) has a residual cost of at least 100%.

References

[1] D. Aadland and V. Kolpin (2004). *Environmental determinants of cost sharing*, Journal of Economic and Behavioral Organization, 53, 495-511, 2004.

[2] E. Anshelevich, A. Dasgupta, E. Tardos and T. Wexler, *Near-optimal network design with selfish agents.* proceedings of the 35th Annual ACM Symposium on the Theory of Computing, 2004.

[3] W. Baumol, J. Panzar and R. Willig, *Contestable markets and the theory of industry structure,* New York: Harcourt Brace Jovanovich, 1982.

[4] Y. Chen and J. Zhang, *Design of price mechanisms for network resource allocation via price of anarchy,* mimeo, New york University, 2005.

[5] M. Cripps and J. Swinkels, *Depth and efficiency of large double auctions,* Econometrica, 74, 1,47-92,2006.

[6] A. Demers, S. Keshaw and S. Shenker, *Analysis and simulation of a fair queuing algorithm,* Internetworking: Research and experience, 1, 3-26, 1990.

[7] J. Green, and J.J. Laffont , *Incentives in public decision making,* Amsterdam: North-Holland, 1979.

[8] R. Johari and J. Tsitsiklis *Efficiency loss in a network resource allocation game,* Mathematics of Operations Research, 29, 3, 407-435, 2004.

[9] R. Johari and J. Tsitsiklis *A scalable network resource allocation mechanism*, mimeo, Stanford and MIT, 2005.

[10] R. Johari, S. Mannor and J. Tsitsiklis *Efficiency loss in a network resource allocation game: the case of elastic supply,* 2004.

[11] E. Koutsoupias and C. Papadimitriou *Worst case equilibria* Proceedings of the 16th Symposium on Theoretical aspects of computer science, 404-413, 1999.

[12] N. Kukushkin, *Best response dynamics in finite games with additive aggregation,* Games and Economic Behavior, 48, 94-110,2004.

[13] S. Littlechild and G. Owen, *A simple expression for the Shapley value in a simple case,* Management Science, 20, 370-2, 1973.

[14] R. P. McAfee, *A dominant strategy double auction,* Journal of Economic Theory, 56, 434-50, 1992.

[15] D. Monderer and L. Shapley, *Potential games,* Games and Economic Behavior, 14, 124-143,1996.

[16] D. Monderer and L. Shapley, *Fictitious play property for games with identical interests*, Journal of Economic Theory, 68, 258-265, 1996.

[17] H. Moulin, *The price of anarchy of serial, average, and incremental cost sharing,* Rice University, 2006.

[18] H.Moulin and S. Shenker, *Strategy-proof sharing of submodular costs: budget balance versus efficiency,* Economic Theory, 18, 3, 511-533, 2001.

[19] R. Musgrave and A. Peacock, *Classics in the theory of public finance,* London: McMillan, 1958.

[20] J.Roemer and J. Silvestre, *The proportional solution for economies with both private and public ownership,* Journal of Economic Theory, 59, 426-444, 1993.

[21] T. Roughgarden *The price of anarchy is independent of the network topology,* STOC, 2002.

[22] T. Roughgarden and E. Tardos *How bad is selfish routing ?* Journal of the ACM, 49, 2, 236-259, 2002.

[23] S. Sanghavi and B. Hajek *Optimal allocation of a divisible good to strategic buyers,* Proceedings of the 43d IEEE conference on Decision and Control, 2004.

[24] M. Satterthwaite and S. Williams, *The optimality of a simple market mechanism,* Econometrica, 70, 1841-1864, 2002.

[25] P.D. Straffin and J.P. Heaney, *Game Theory and the Tennessee Valley Authority,* International Journal of Game Theory,10, 1,35-43,1981.

[26] M. Walker, *On the non-existence of dominant strategy mechanisms for making optimal public decisions,* Econometrica, 48, 1521-40, 1980.

[27] S. Yang and B. Hajek *Revenue and stability of a mechanism for efficient allocation of a divisible good,* mimeo, University of Illinois at Urbana Champaign, 2005.