Part 4: Topics in quenching

Recap: What do we mean by quenching?

Quenching is the presumed process that suppresses star-formation in some galaxies to a level 1-2 dex, or more, relative to the level in "Main Sequence" galaxies of the same mass. It transforms "active" star-forming galaxies that are forming stars at a cosmologically significant rate, $sSFR \geq \tau_H^{-1}$, into "passive" galaxies that are not forming stars at a significant rate $sSFR \ll \tau_H^{-1}$.   Most stellar mass is in passive galaxies.

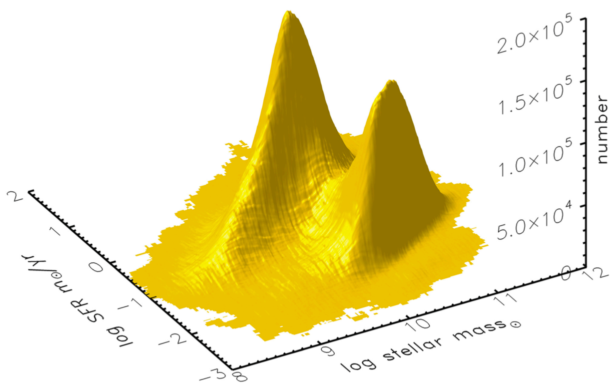

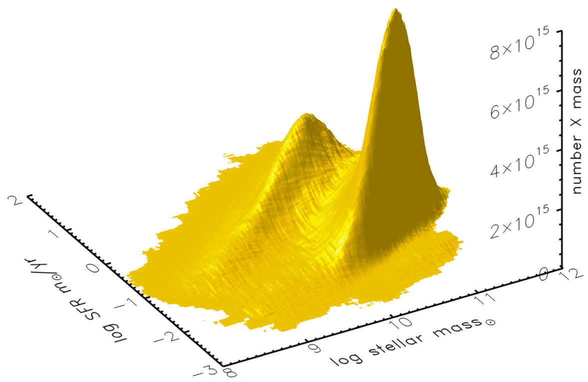

We also saw how the cessation of significant star-formation in quenched galaxies produces the roll-over in the m_{star}/m_{DM} ratio at high masses, and contributes to the decline in the SFRD.

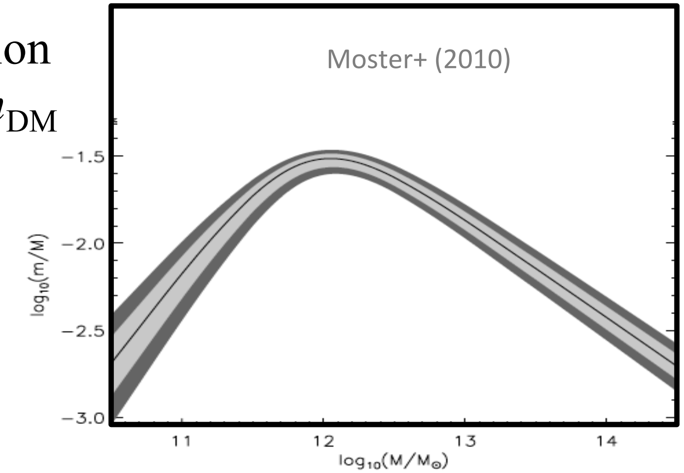

In order to quench, we presumably need to do one or more of the following:

- Stop gas inflow from the halo
- Remove gas from the galaxy
- Stop making stars out of gas (very low ε)

Mass and environment (Peng et al 2010, 2012)

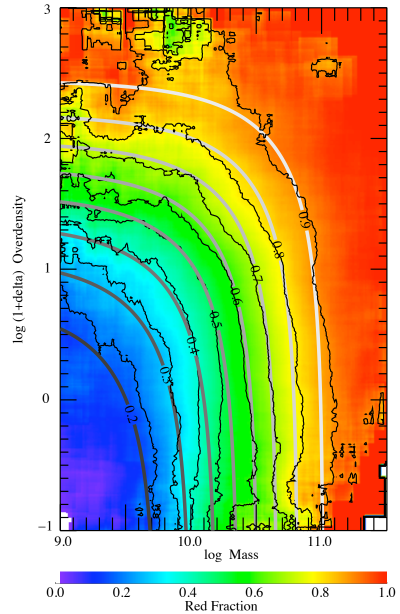

We saw how the fraction of objects still making stars is separable in stellar mass and an environment variable (ρ = local density of galaxies), suggesting two channels of quenching, mass quenching and environment-quenching.

$$f_{blue}(m,\rho) = \left(1-\varepsilon_m(m)\right)\times\left(1-\varepsilon_\rho(\rho)\right)$$

Environment-quenched galaxies are almost all *satellite* galaxies. We can then define the "satellite quenching efficiency" as the chance that a satellite is quenched because it is a satellite, i.e. compared with what it would have been as a central of the same mass.

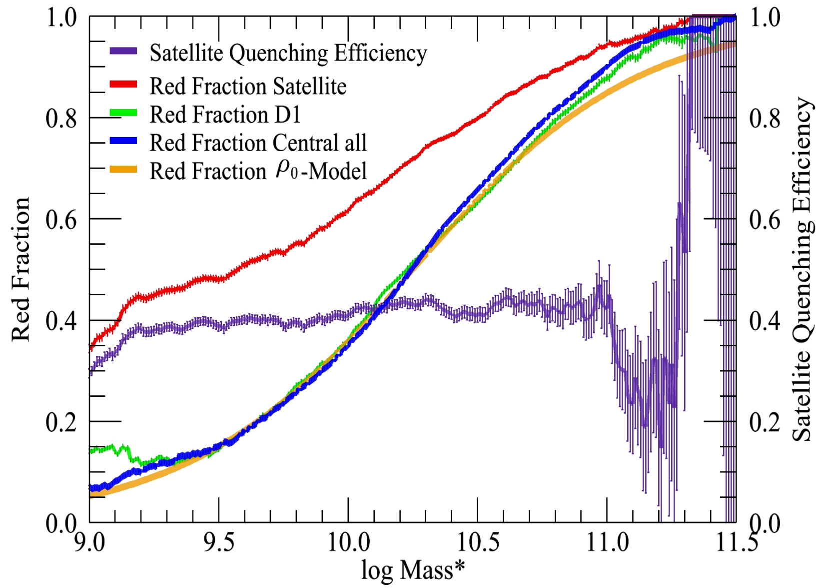

$$\varepsilon_{sat}(m_*, p_1, p_2 \ldots) = \frac{f_{q|sat}(m_*, p_1, p_2 \ldots) - f_{q|cen}(m_*)}{f_{q|cen}(m_*)}$$

Not surprisingly, we find that ε_{sat} is strikingly independent of stellar mass. We'll see what else it depends on later.

Satellite quenching is due to the "environment", due to a galaxy being a satellite.

What about mass-quenching which affects centrals at the bottom of their potential well? Is this driven by internal processes in the galaxy**, or external ones in the halo or even further afield, or a combination of the two?

** *unfortunately, stellar mass is correlated with many other things, including galactic structure, mass of the central BH, DM halo mass, etc., making it not easy to distinguish.*

Four <u>very diverse</u> topics for today's lecture

- *Physics*: The basic problem of (mass-) quenching: why do we need it?
- *Data interpretation*: Surface mass-density as a driver of quenching? The difficulties of establishing causality and progenitor bias.
- *Physics*: Mechanisms for satellite quenching
- *Philosophy(?)*: Conformity, knowledge and ignorance

Then, next week three more.

- *Reverse Engineering*: Supermassive Black Holes and Active Galactic Nuclei
- *Physics:* Energy injection from black holes: observational evidence
- *Scientific approach*: Is quenching even a real concept? The importance of your paradigm.

Topic 1:
Why do we need to actively quench centrals?

Why do we need quenching of centrals?

Much earlier in the course we talked about cooling in haloes and I argues that higher mass haloes are less able to cool. The relation between virial temperature and mass of haloes is as follows:

$$T_{vir} \sim 10^{6}\left(\frac{M}{10^{12} M_{sun}}\right)^{2/3}\left(\frac{1+z}{3}\right)\mathrm{K} \sim 36\left(\frac{V_{vir}}{\mathrm{kms^{-1}}}\right)^{2}$$

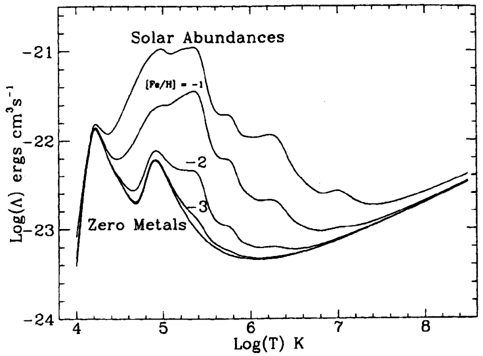

So, 10^{12} $M_\odot$ corresponds to ~ 3×10^{5} K, close to the drop in the cooling curve. Is this all there is to it?

It's part of the story, but not the whole of it.

Gas entering the halo is shocked, i.e. the kinetic energy of bulk inflow is thermalised (converted to random motions) to the virial temperature as part of "virialization".

Gas heated to $T_{\mathrm{vir}} < 10^{5.5}$ K rapidly cools down to ~ $10^{4.5}$ K post-shock, and can flow down to the galaxy, feeding star-formation, as we have seen.

In haloes above a threshold of around 10^{12} $M_\odot$ the shocked gas joins a reservoir of hot gas at $T_{\mathrm{vir}} > 10^{5.5}$ K. What happens then?

It is not as simple as this: there will always be parts of the halo which are cooling effectively because of their high density. Indeed, we see the hot gas radiating strongly in X-rays (esp. in the centers of the haloes), so the gas is demonstrably losing energy (cooling), and should be sinking. There really ought to be some gas flow onto the central object.

(*This was realised in mid-1970's when clusters of galaxies were found to be strong X-ray sources*).

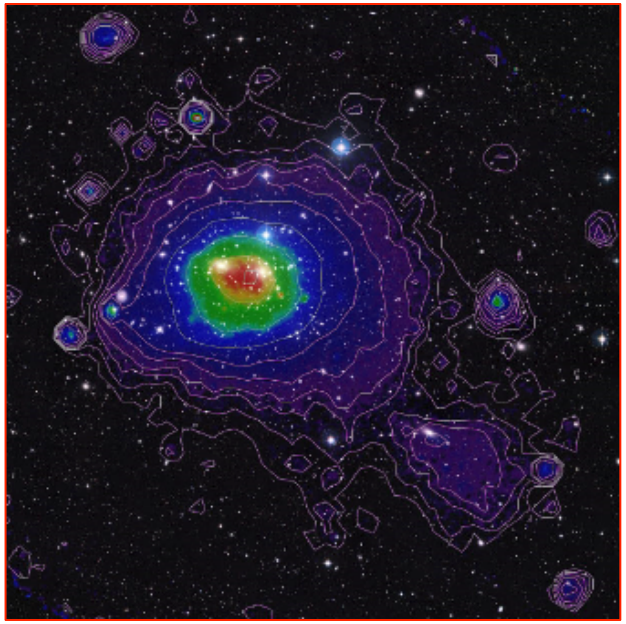

We had before that the cooling time is just given by the density and temperature at a particular radius.

$$t_{cool} \sim \frac{3}{2} \frac{\bar{\mu} m_p \, kT(r)}{\rho_{gas}(r) \, \Lambda(T)}$$

We'll assume that the hot gas has a simple isothermal distribution with the virial temperature, so the density goes as

$$\rho_{gas}(r) \sim \frac{m_{gas}}{4\pi R_{vir} r^2}$$

Now, what timescale to equate to t_{cool}?
Age of the Universe? Better agreement with simulations comes from "dynamical time" $t_{dyn} \sim R_{vir}/V_{vir}$, i.e. the halo crossing time. This, it turns out, is a more or less constant fraction of the age of the Universe.

$$\frac{R_{vir}}{V_{vir}} \sim 0.1\, H^{-1} \sim 0.1 \tau_H$$

From Croton et al 2006

There will be some "cooling radius" r_{cool} at which $t_{cool} = t_{dyn}$

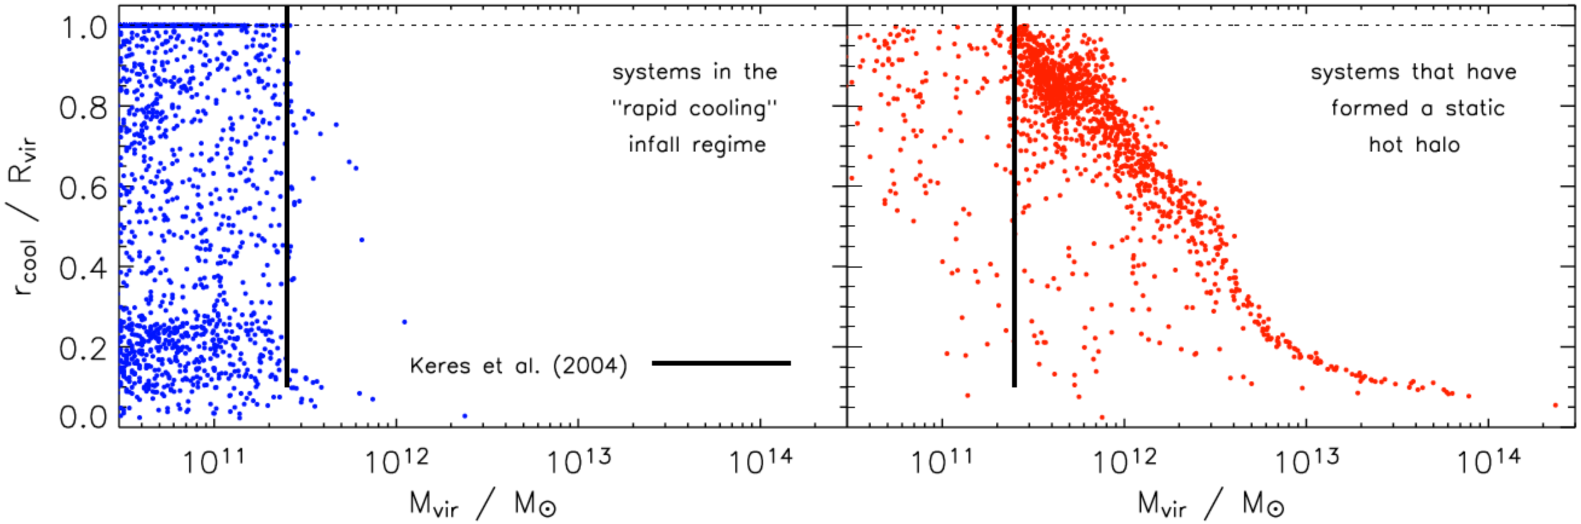

We can now get a rate at which gas should be deposited on the galaxy by seeing how fast the cooling radius changes with time. Bottom line: Even if the cooling radius is small, we should get a significant deposition rate of gas.

$$\dot{m}_{cool} \sim 4\pi\, \rho_{gas}(r)\, r_{cool}^2\, \dot{r}_{cool}$$
$$\sim 5\, m_{gas} \frac{r_{cool}}{R_{vir}} \frac{1}{\tau_H}$$

$$\dot{m}_{cool} \sim 5\, m_{gas} \frac{r_{cool}}{R_{vir}} \frac{1}{\tau_H}$$

The problem is therefore how to stop the gas in haloes cooling onto the galaxy! *Important point: the problem of quenching is therefore not (only) how to quench star-formation in a galaxy, the real problem is* <u>*how to keep it quenched for the rest of time.*</u>

How much energy would we need to inject to prevent this mass inflow? Basically, we'd need to heat this incoming mass back up to the virial temperature.

$$\text{Heating rate} \sim \frac{1}{2} \dot{m}_{cool} V_{vir}^2$$

Supernova heating? Probably not. We saw earlier that if we form some mass of stars, the resulting SN can in principle (at maximum efficiency) "heat" the same mass of gas to $v \sim 500$ kms^{-1}, well below the $V_{\rm vir}$ of more massive haloes, and anyway requiring substantial star-formation at a level comparable to $dm_{\rm cool}/dt$, which is not seen.

$E_{\rm SN} \sim 4\times10^{48}$ erg $M_\odot^{-1}$ was due, ultimately to the energy release when a small fraction of the mass collapses to neutron star densities, inefficiently coupled to matter (most energy in a neutrino pulse).

What about black holes? The energy that can in principle be extracted when mass *m* accretes onto a black hole is potentially <u>enormously</u> higher.

$$E_{\rm BH} \sim 0.1\, mc^2 \sim 2 \times 10^{53} \text{ erg } M_\odot^{-1}$$

It is widely accepted that this is (frankly) the only plausible energy source, even though the details are very poorly understood (see next week).

Topic 2:
Surface density and the perils of interpretation:
A cautionary tale of interpretation

Observationally there are certainly strong links between quenching and surface density and/or other manifestations of galactic structure

The fraction of quenched galaxies is closely correlated with the stellar surface mass density Σ (either average $\Sigma_e \sim m_{star}/R_e^{-2}$ or simply $\Sigma_{1\ kpc}$). Σ is clearly a *much* better predictor than mass of whether you are quenched or not, no doubt about it.

Also similar effects with velocity dispersion (Smith+ 2009, Wake+ 2012), Sersic indices (Blanton+ 2003, Wuyts+ 2011), bulge mass (Bluck+2015)

Structure could produce quenching directly, e.g. via disk stability (Martig+2009, Genzel+ 2014) or indirectly (e.g. via black hole mass following m_{bulge} etc)

So, a lot of interest in the role of Σ (i.e. structure) in quenching.

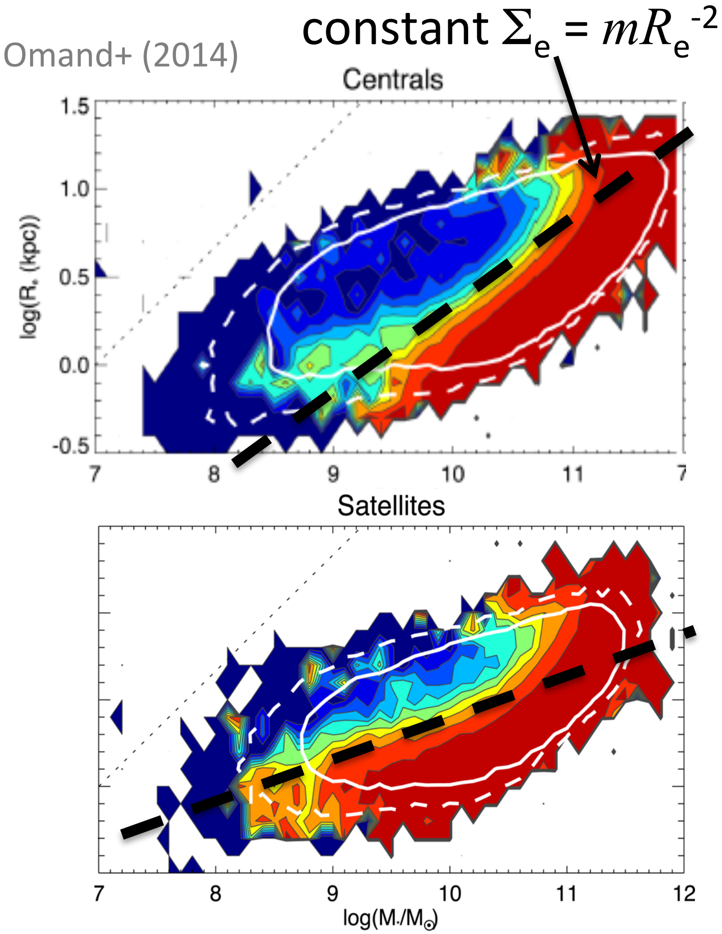

"… in any case it reinforces *the fact* that is not correct to assume that galaxy mass is the driving parameter (e.g. Peng et al. 2010) ".

"… it implies that satellite quenching *must* be accompanied by a change in structure. "

Omand et al 2014

But reasons (for me) to be suspicious….

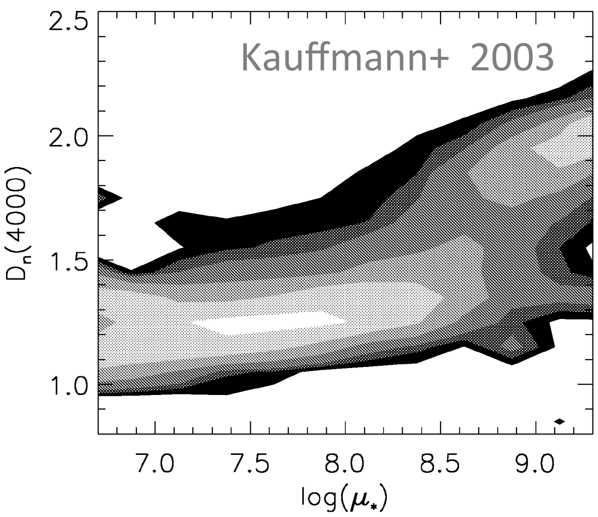

The 2014 Omand et al plot was an illustration of a result known for some time: as Σ (here written as μ) increases, there is an abrupt transition from star-forming (low "D4000") to passive (high "D4000") galaxies.

Furthermore, there was evidence that this threshold Σ was higher at high z, increasing roughly as $(1+z)^2$. i.e. star-forming galaxies have higher Σ at high z than do star-forming galaxies of the same mass today.

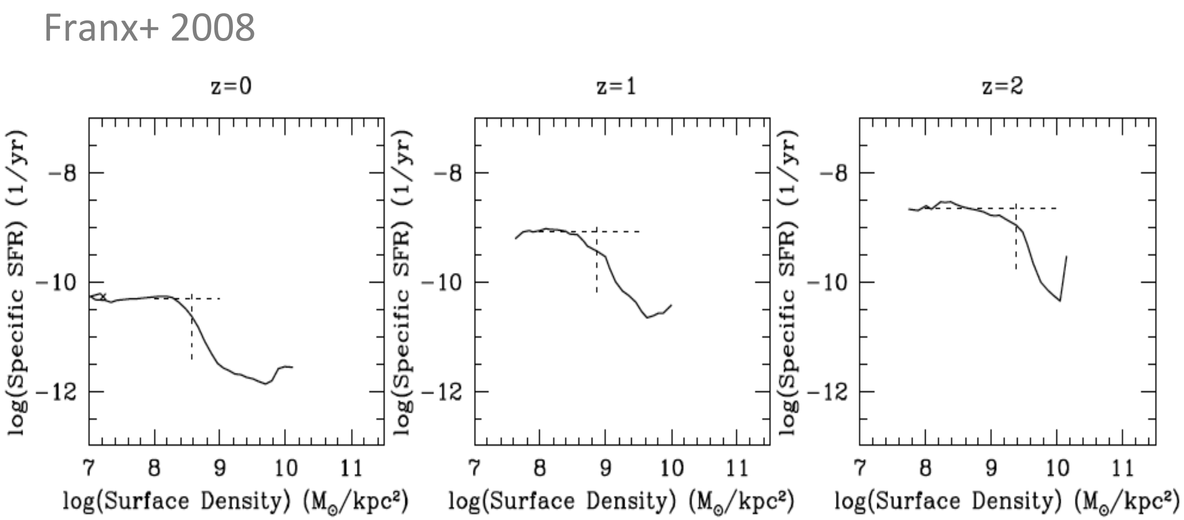

This should not surprise you: we know that DM haloes of the same mass are $(1+z)$ smaller at high z (so that the density scales as the cosmic density), so with constant collapse factor for the baryons, we might expect Σ to be higher as $(1+z)^2$.

Indeed, such size “evolution”** is seen

** N.B. This is an evolution of the *population* of galaxies at the same mass, not of individual objects, which are continuously increasing in mass (and size).

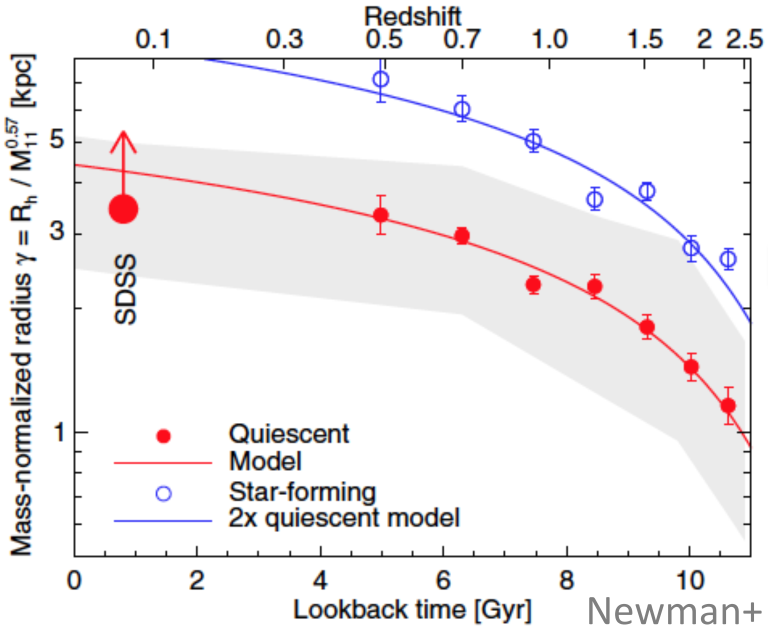

Newman+ 2012

also Buitrago+ 2008, Mosleh+ 2014+….

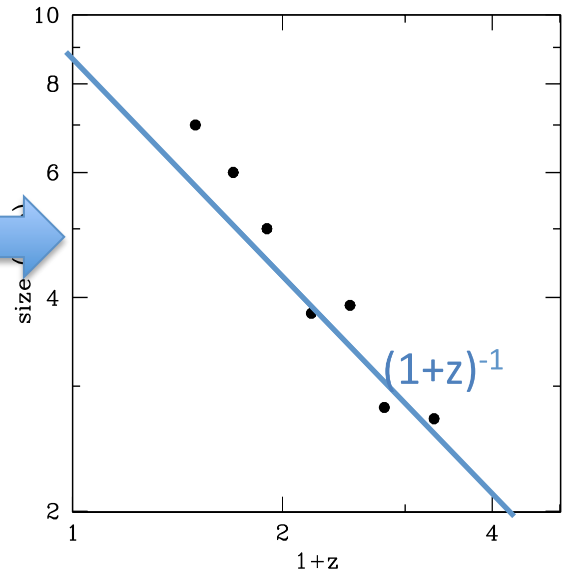

Also, at a given stellar mass, passive galaxies are systematically smaller than star-forming ones.

At a given stellar mass, the size of star-forming galaxies indeed scales roughly as $(1+z)^{-1}$, i.e. Σ as $(1+z)^{-2}$

Leading to an obvious question: Are quenched galaxies small and dense because density played a role in quenching them, or because they quenched (i.e. stopped evolving), for some <u>other</u> reason, at high redshift, when all galaxies were denser?

A simple toy-model (Lilly & Carollo 2016)

Addition of new stars

Galaxies form stars at a rate given by their existing mass and the Main Sequence sSFR(m,z) and add these in an exponential disk with scale length given by:

$$h_{SF} = 5\left(\frac{m}{3\times 10^{10}}\right)^{1/3}(1+z)^{-1}\ \mathrm{kpc}$$

Individual objects: Inside-out growth due to both m and z terms

Quenching

Galaxies then probabilistically cease star-formation, instantaneously and throughout and with no rearrangement of mass, according to the Peng+ (2010) prescription for "mass-quenching", which may be written as

$$Pdm = \frac{dm}{M*}$$

plus satellites have an additional probability of satellite-quenching Peng+ (2010,2012) given by:

$$P_{sat}dm = \varepsilon_{sat}\frac{dm}{m}(1+z)^{-1}$$

- absolutely <u>no</u> dependence of quenching on size or surface density.
- only difference between centrals and satellites is in the <u>mass dependence</u> of quenching.
- Absolutely no structural mass rearrangement on quenching.

Recovery of size evolution and offset SF-passive

Passive galaxies are typically factor of two smaller than star-forming ones (as observed)

This is due to both differential fading (F) and “progenitor bias**” (P), both linked to smaller sizes in the past.

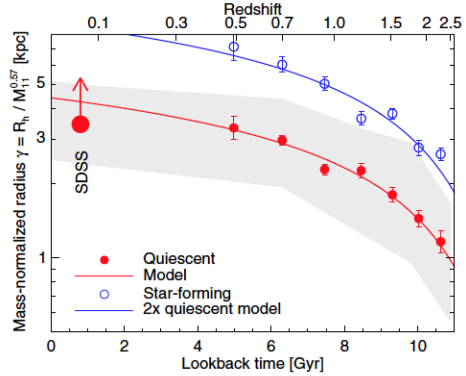

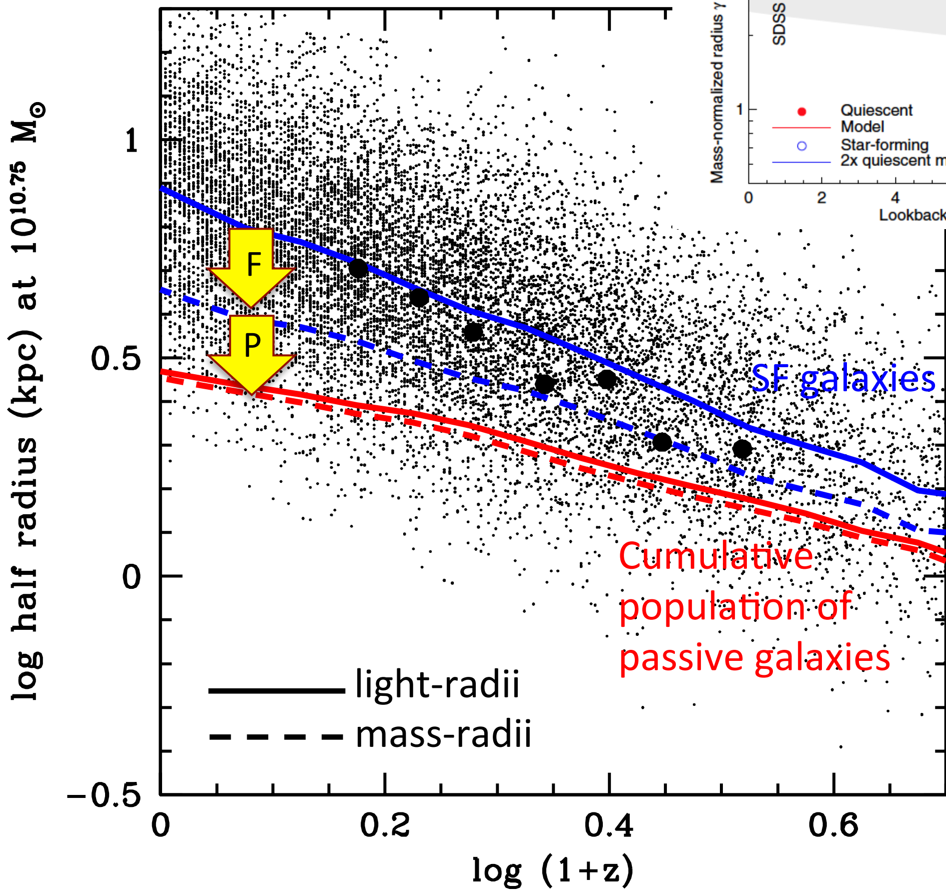

*** what is “progenitor bias”?*
The properties of a population (e.g. half-light sizes of passive galaxies) changes with epoch because of changes (e.g. additions) to the population rather than to changes in the properties of individual members. It is likely whenever the new members of a population are likely to have different properties to existing members.

The quenching stellar density thresholds follow from size evolution…

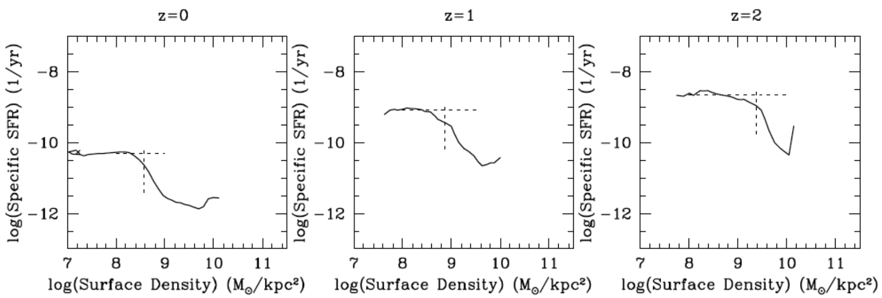

Franx et al (2008)

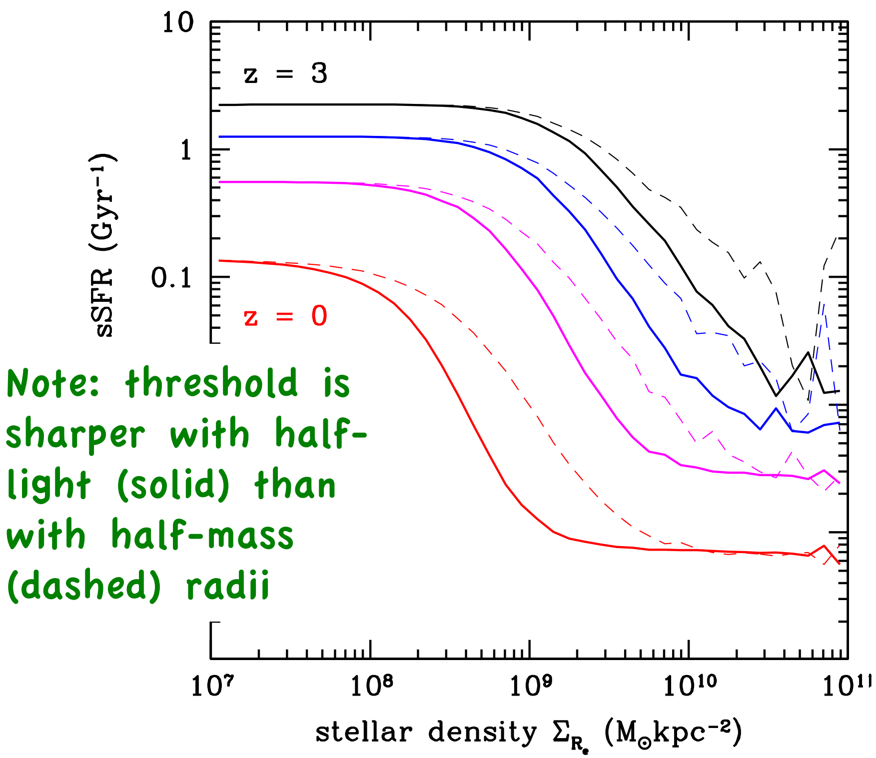

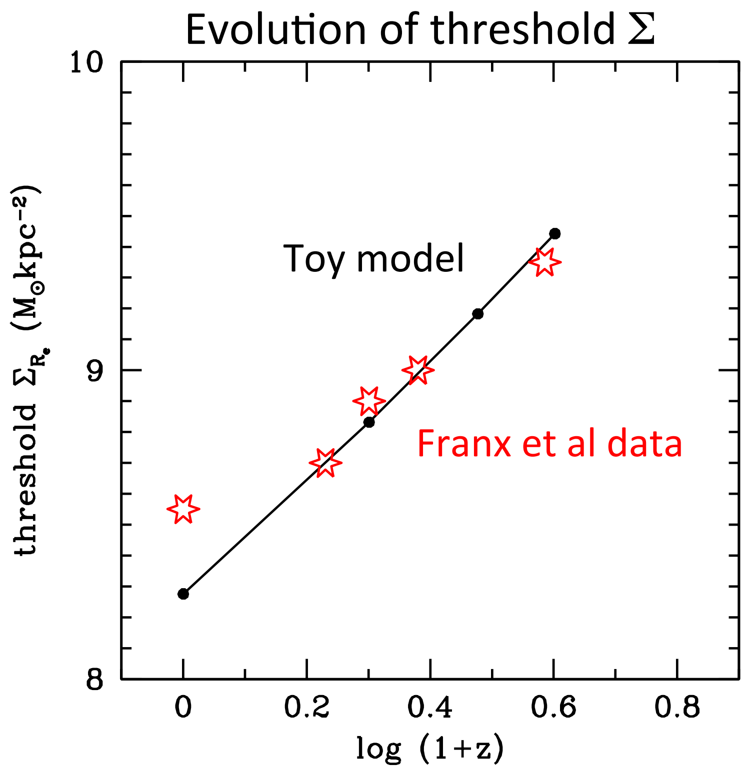

... and the apparent role of stellar density in f_Q arises naturally

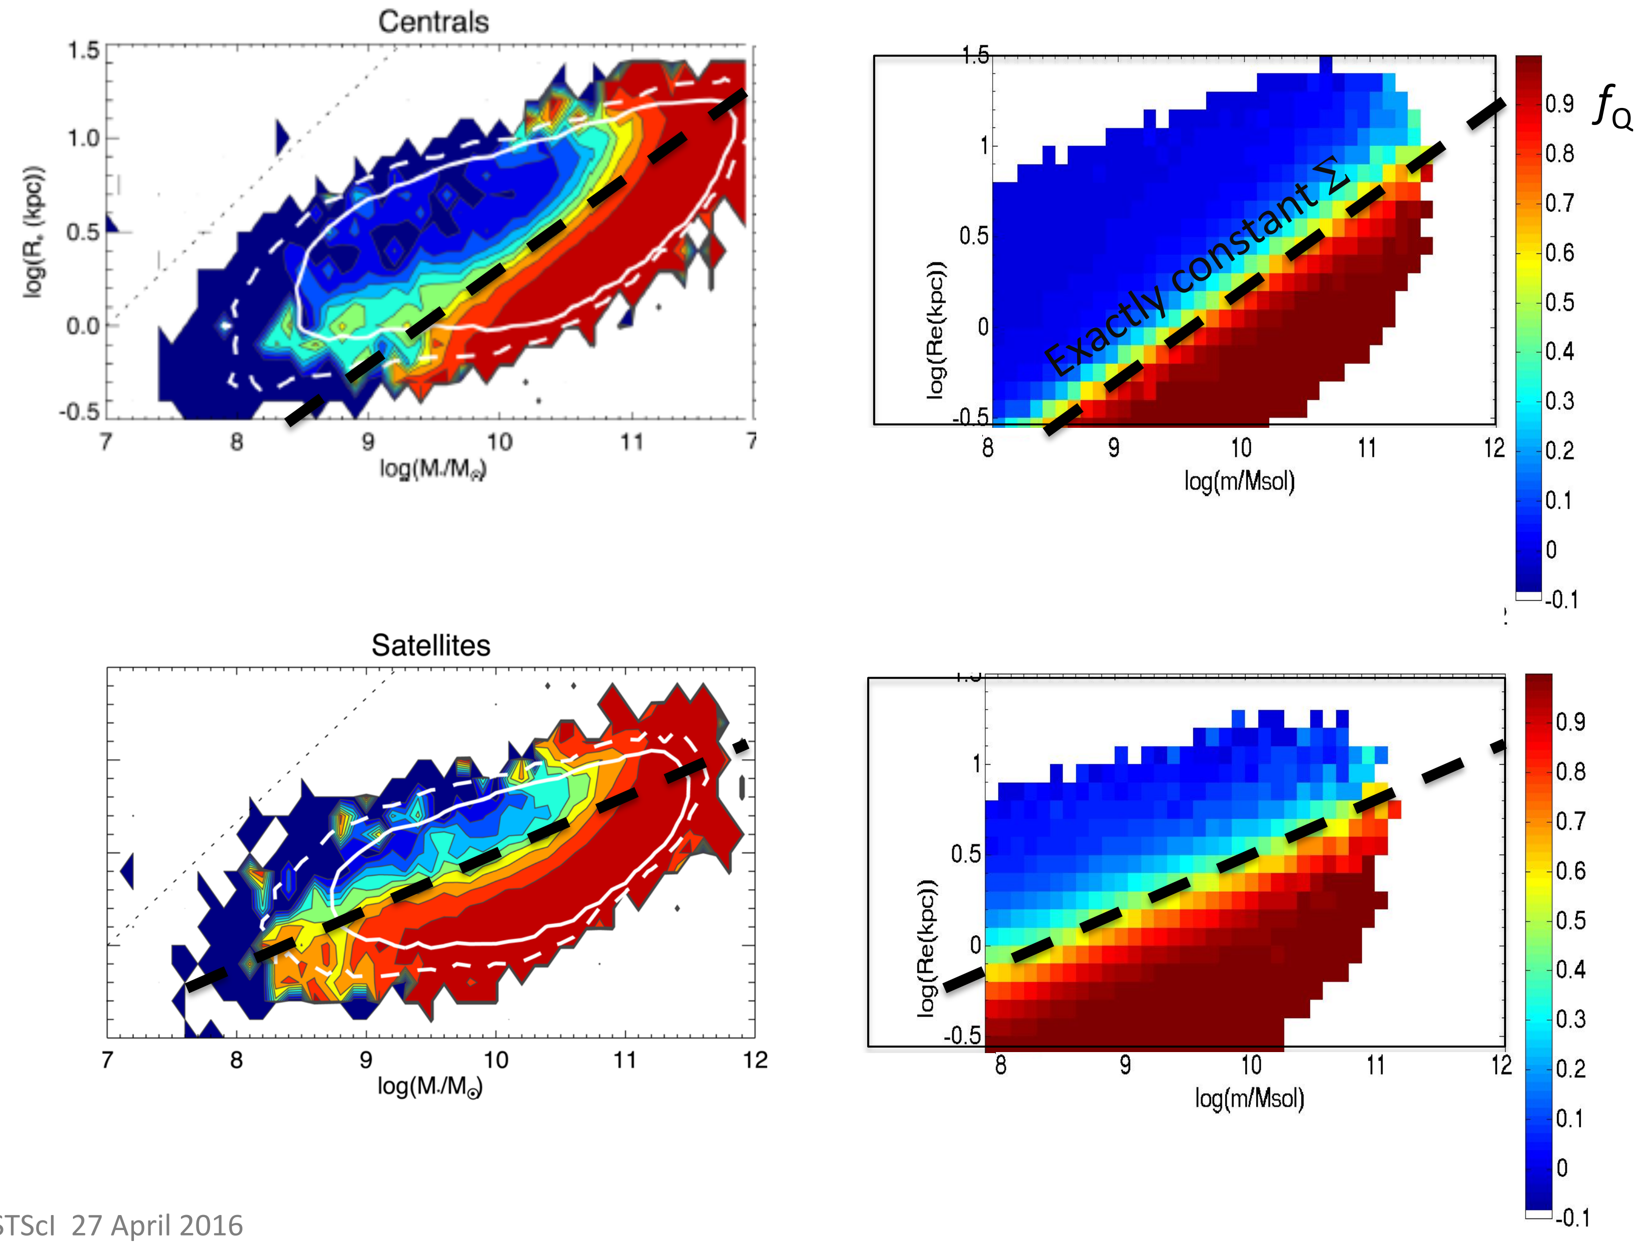

... and the apparent role of stellar density in f_Q arises naturally

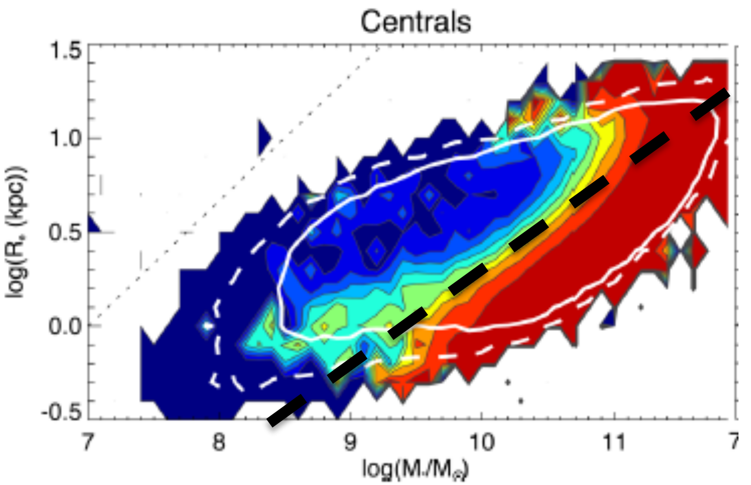

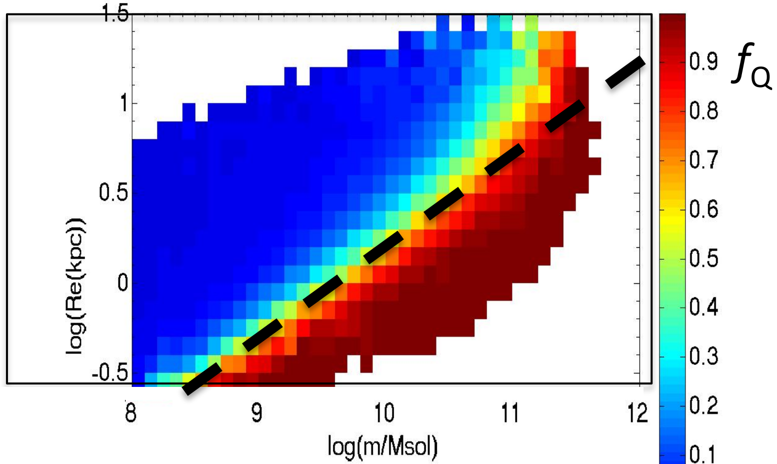

With a very naïve effect of mergers on quenched galaxies: Δm = +0.2 dex, ΔR_e = +0.2 dex

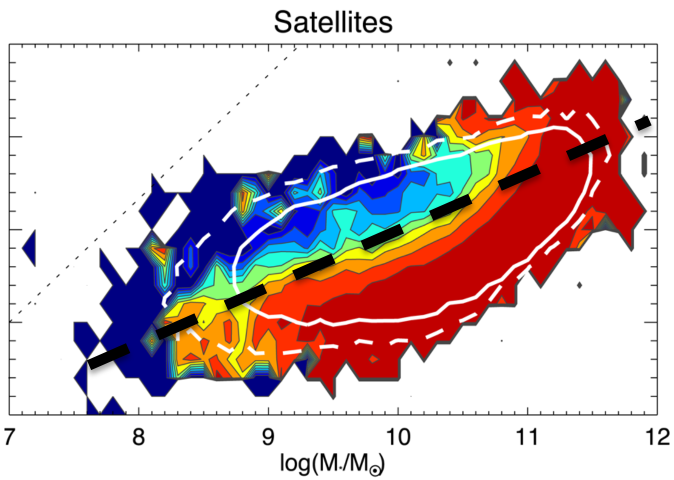

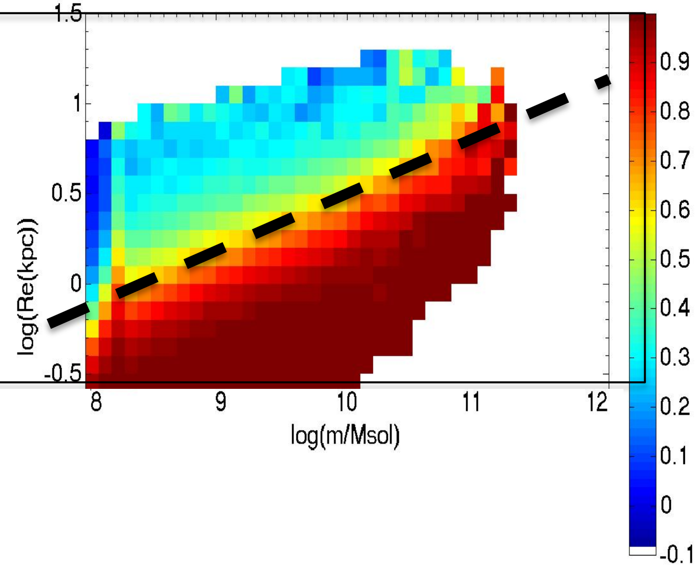

Where Σ_{thresh} comes from and why satellites are different than centrals

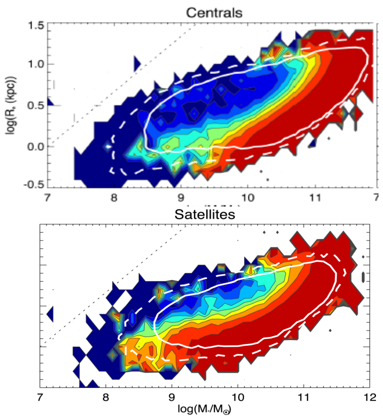

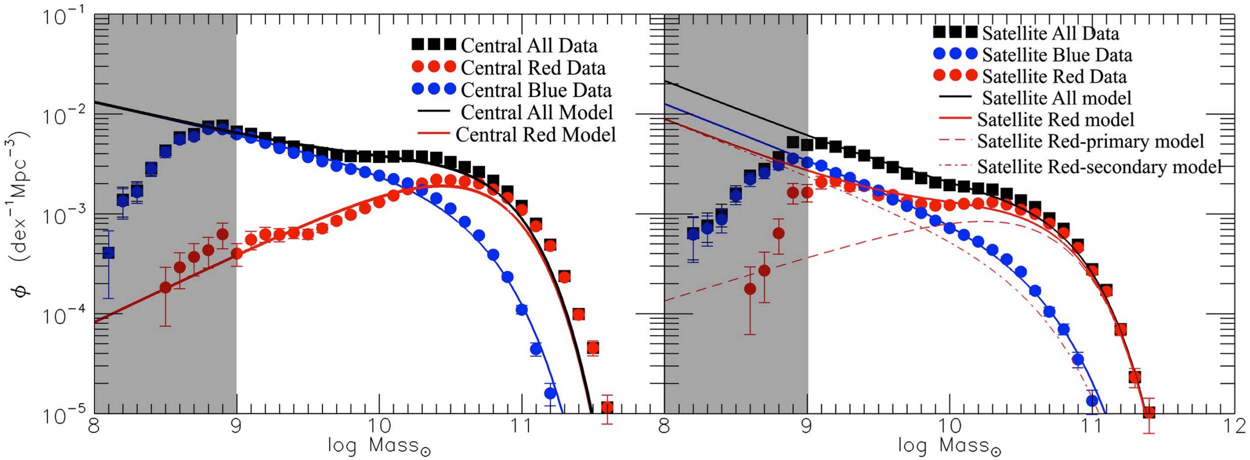

average star-forming $R_e(m)$
average all R_e(m)
average quenched $R_e(m)$

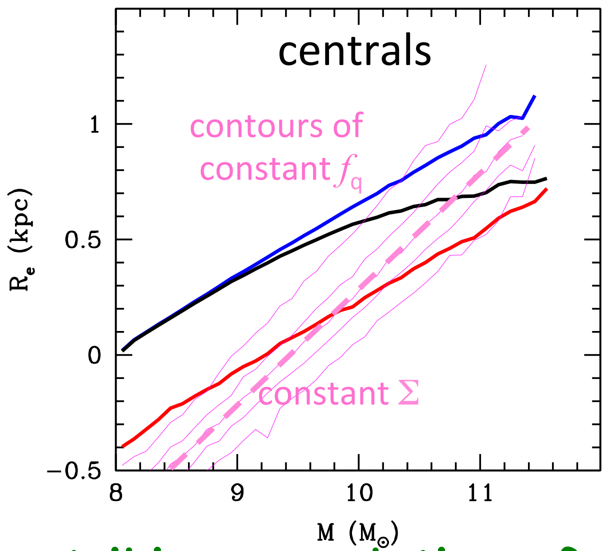

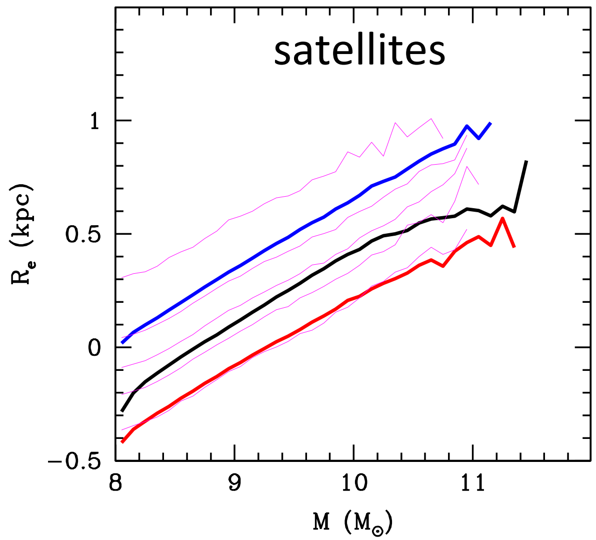

Conclusion: The striking correlation of quenched fraction and surface mass density Σ is almost certainly (in my view) a coincidental result of quite unrelated quenching processes

Where Σ_{thresh} comes from and why satellites are different than centrals

Should you believe this? Isn't it based on a naively simple toy-model?

Actually it is ultimately based on three observational facts:

- Factor of roughly two offset in $R_e(m)$ between star-forming and passive galaxies
- Differences in mass functions of star-forming and passive, and between centrals and satellite galaxies
- The roughly $R_e \sim m^{1/3}$ mass-size relation at given epoch.

The toy-model just reproduces and unifies these three into a single coherent scheme based on it's two main inputs:

$$h_{SF} = 5\left(\frac{m}{3\times10^{10}}\right)^{1/3}(1+z)^{-1}\ \text{kpc}$$

Evolving size-mass relation

$$Pdm = \frac{dm}{M^*} \quad \text{and} \quad P_{sat}dm = \varepsilon_{sat}\frac{dm}{m}(1+z)^{-1}$$

Peng+ mass-dependent quenching

The difficulty of establishing the direction of causality in a passive science

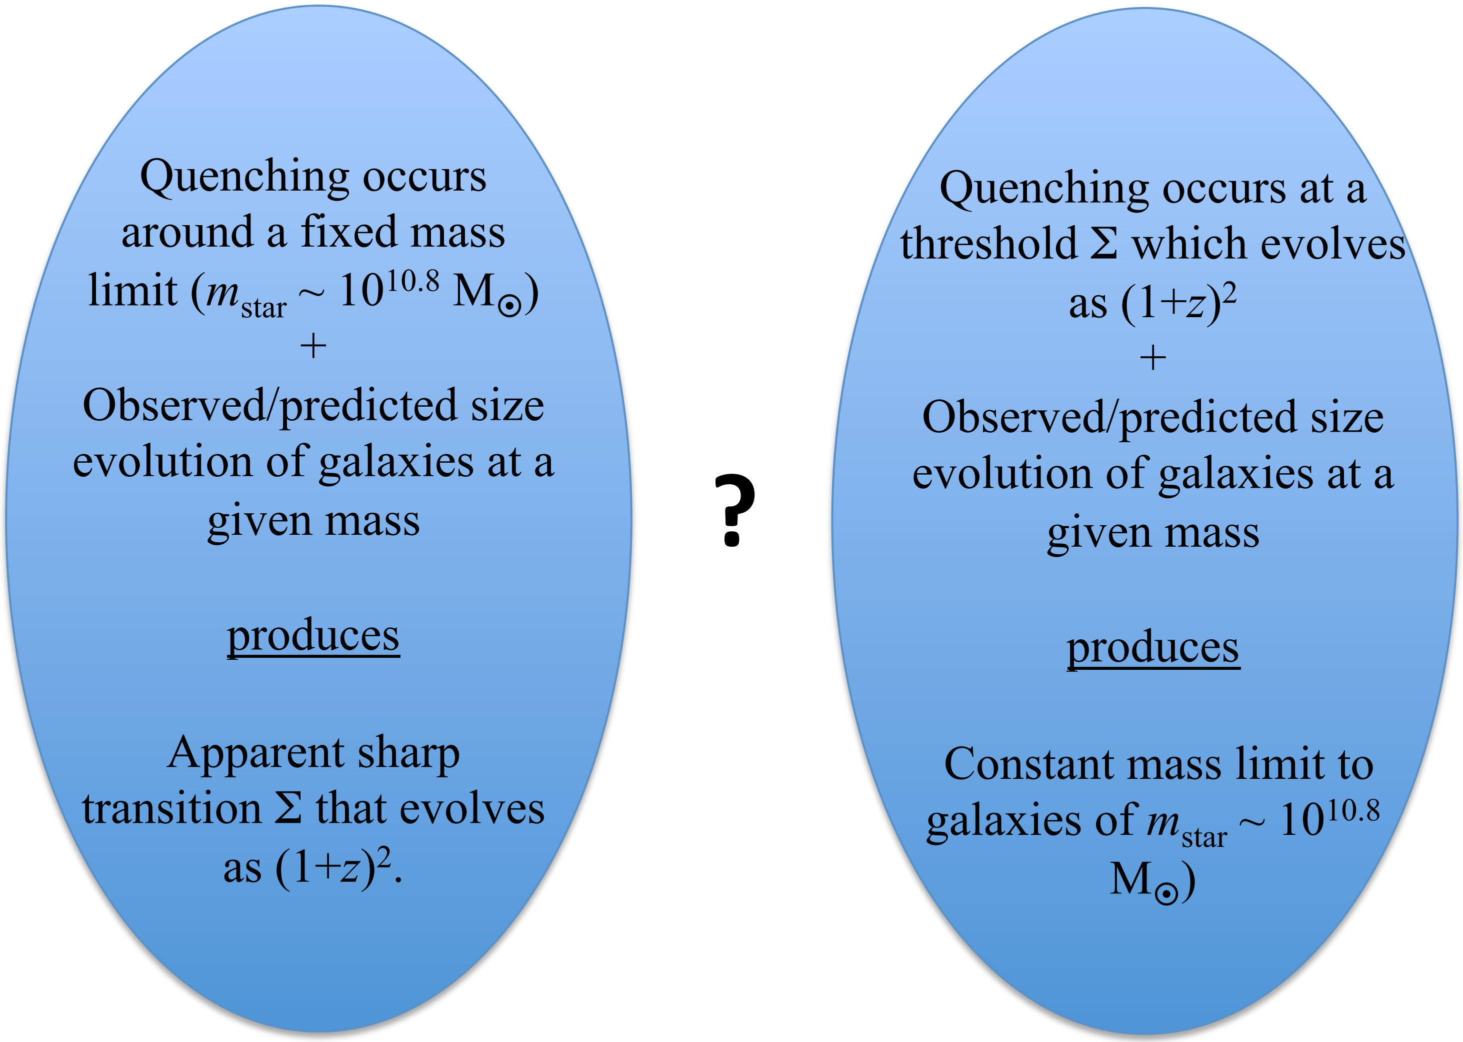

Topic 3:
Satellite quenching

A. Tidal stripping of satellites

Suppose we have a satellite galaxy orbiting in the potential of another more massive galaxy. Tidal effects from the gravitational field of the larger galaxy can strip material from the smaller one. Tidal effects arise from the gradient of the gravtitational field.

The following simplified analysis shows how this works. For analytic simplicity we we imaging the two galaxies to be point masses, of mass $M_{host} >> M_{sat}$ respectively, separated by d, and that the satellite is on a circular orbit with angular frequency ω. We then consider a test-particle of mass m in the satellite on a circular orbit of radius r.

This is the circular restricted 3-body problem. If we look at the frame that is co-rotating with the satellite around the host, we find the two Lagrangian points L1 and L2.

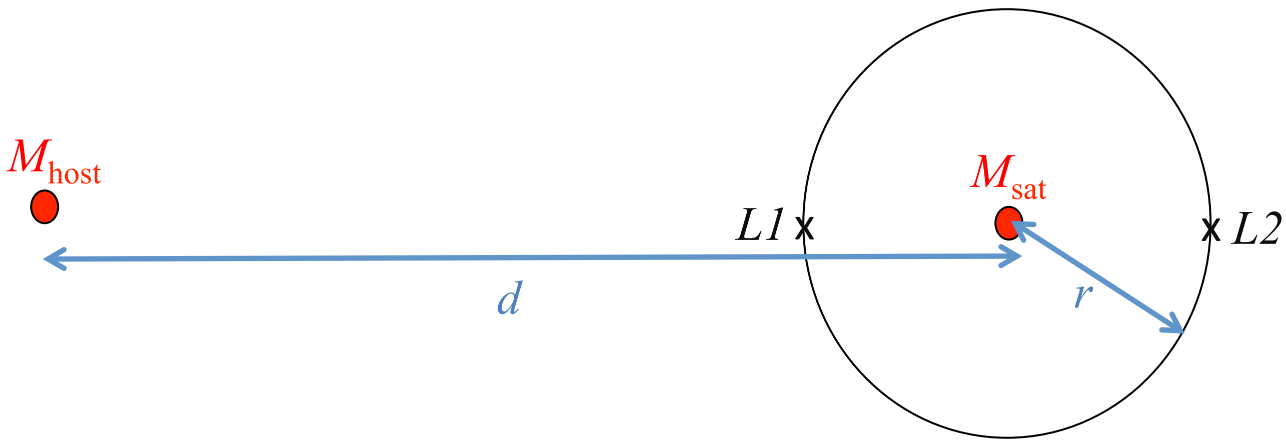

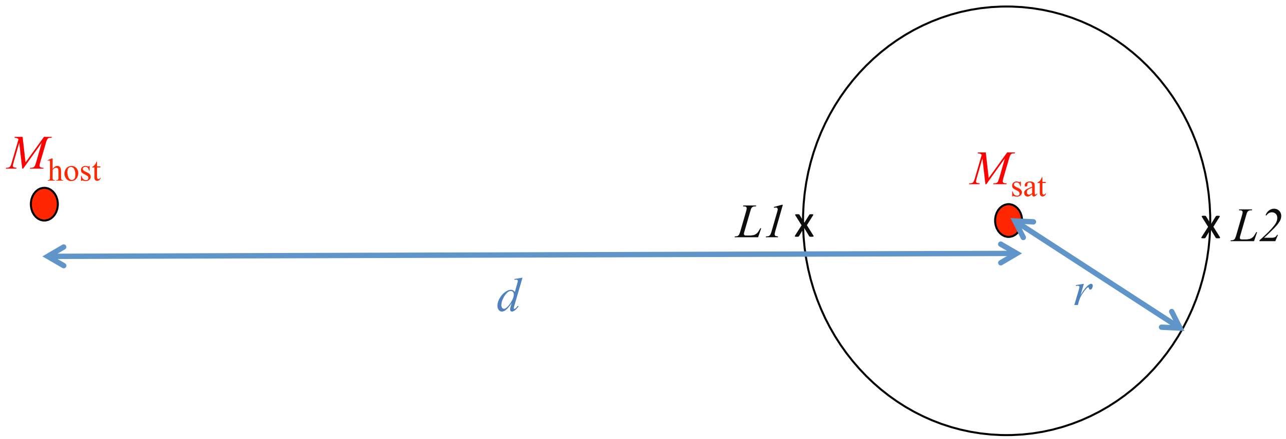

L1 and L2 are defined such that the total gravitational force on a particle at these positions causes motion around the host with exactly the same angular speed ω as the satellite, so that the particle stays *fixed* in the co-rotating frame. Particles at greater distance from the satellite than will not be bound and will move away (be "tidally stripped") from the satellite. For $M_{host} >> M_{sat}$, L1 and L2 are equidistant (at distance r_t) from M_{sat},

For the satellite galaxy:

$$\frac{GM_{host}M_{sat}}{d^2} = M_{sat}d\omega^2$$

For test-mass in the satellite (the + will refer to the L2 position and the – to L1)

$$\frac{GM_{host}m}{(d \pm r_t)^2} \pm \frac{GM_{sat}m}{r_t^2} = m(d \pm r_t)\omega^2$$

For $M_{host} >> M_{sat}$, which implies $d >> r$, then it is easy to solve these to yield the r beyond which material is stripped.

$$r_t \sim \left(\frac{M_{sat}}{3M_{host}}\right)^{1/3} d$$

$$r \sim \left(\frac{M_{sat}}{3M_{host}} \right)^{1/3} d$$

The expression can be easily re-written for (spherically symmetric) extended mass distributions in terms of the mean densities ρ_{host} and ρ_{sat} that are enclosed within *d* and *r* respectively. Material with lower density than $<\rho>_{sat,t}$ will be stripped from the satellite.

$$\langle \rho \rangle_{sat,t} \sim 3 \langle \rho \rangle_{host}$$

The material stripped from the satellite forms a stream along the orbit (behind *and* in front of the satellite). These tidal streams are often observed, e.g. in our own halo.

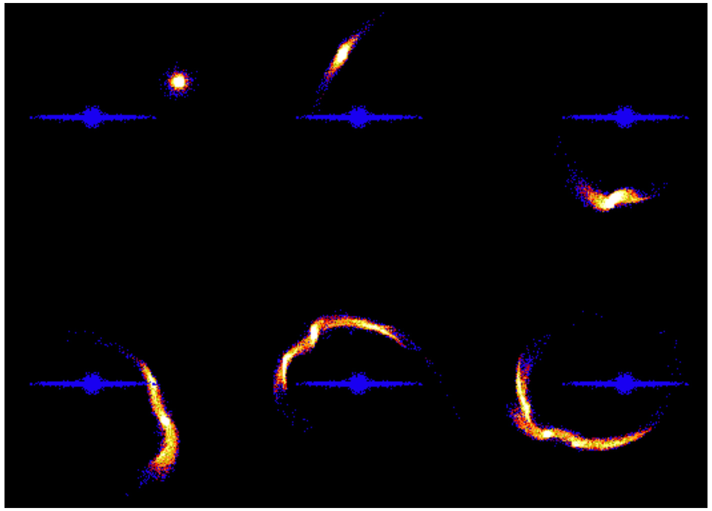

It is easy for satellite galaxies to lose their own (lower density) DM haloes, and not difficult if they pass close enough to the host, to also lose much of their stars and gas.

B. Ram-pressure stripping of satellites

For the gas, there is an additional affect of ram-pressure stripping from diffuse gas in the host galaxy halo (the intracluster medium, ICM). Stars are immune to this because of their tiny cross section).

Gas in a galaxy moving with velocity v through a gaseous medium of density ρ_{ICM} experiences a ram pressure P

$$P = \rho_{ICM} v^2$$

If this exceeds the gravitational force holding the gas to the satellite, then the gas will be stripped away. The gravitational force will be given by the circular velocity v_{sat} of the satellite system. The condition for gas to be lost can then be written approximately as

$$\rho_{ICM} v^2 > \rho_{ISM} v_{sat}^2$$

In a big cluster, it is easy for $v > 1000$ kms-1c.f. $v_{sat} < 200$ kms^{-1}.

Ram pressure stripping in action: ESO 137-001 in Norman cluster

C. Dynamical friction

One other effect is of interest. An object moving of mass M through a diffuse medium of density ρ (and Maxwellian velocity dispersion σ) experiences a drag force because of its own gravitation effect on the particles in the medium, causing the object to lose energy (and spiral towards the center of the potential well.

In the frame of the moving object:

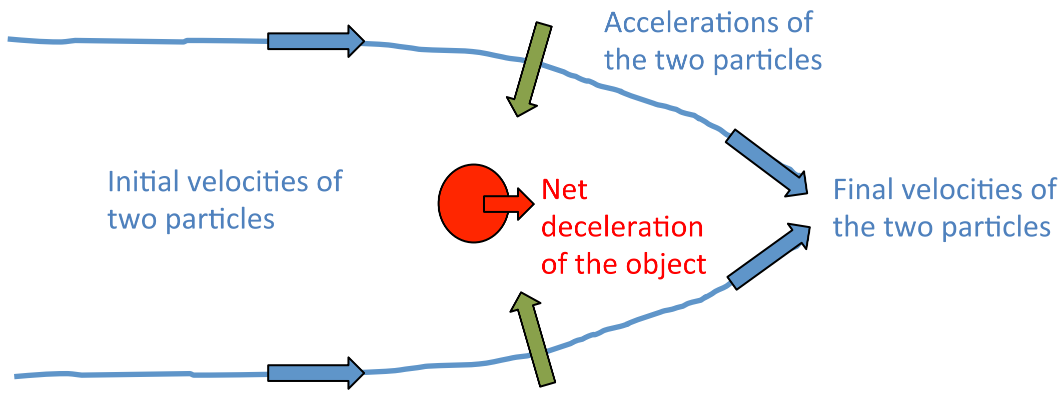

It is not trivial to derive this, so let me state the result (from Chandrasekhar 1943)

$$F = -4\pi \frac{G^2 \rho M^2}{v^2} \ln\Lambda \, F\left(\frac{v}{\sqrt{2}\sigma}\right)$$

ln Λ and F(v/σ) are quantities of order unity….

Note that the force goes as M^2, and does not depend on individual masses of the particles in the medium, only their density More massive objects decelerate faster (this is because the effect is caused by gravity). If we now calculate a timescale for this effect becoming significant, as v divided by dv/dt, we get, as a rough estimate

$$t_{fric} \sim 3\left(\frac{\sigma}{200 kms^{-1}}\right)\left(\frac{r}{100kpc}\right)^2\left(\frac{M_{sat}}{10^{11} M_O}\right)^{-1} Gyr$$

How does $\varepsilon_{sat}(\ldots)$ vary with different parameters?

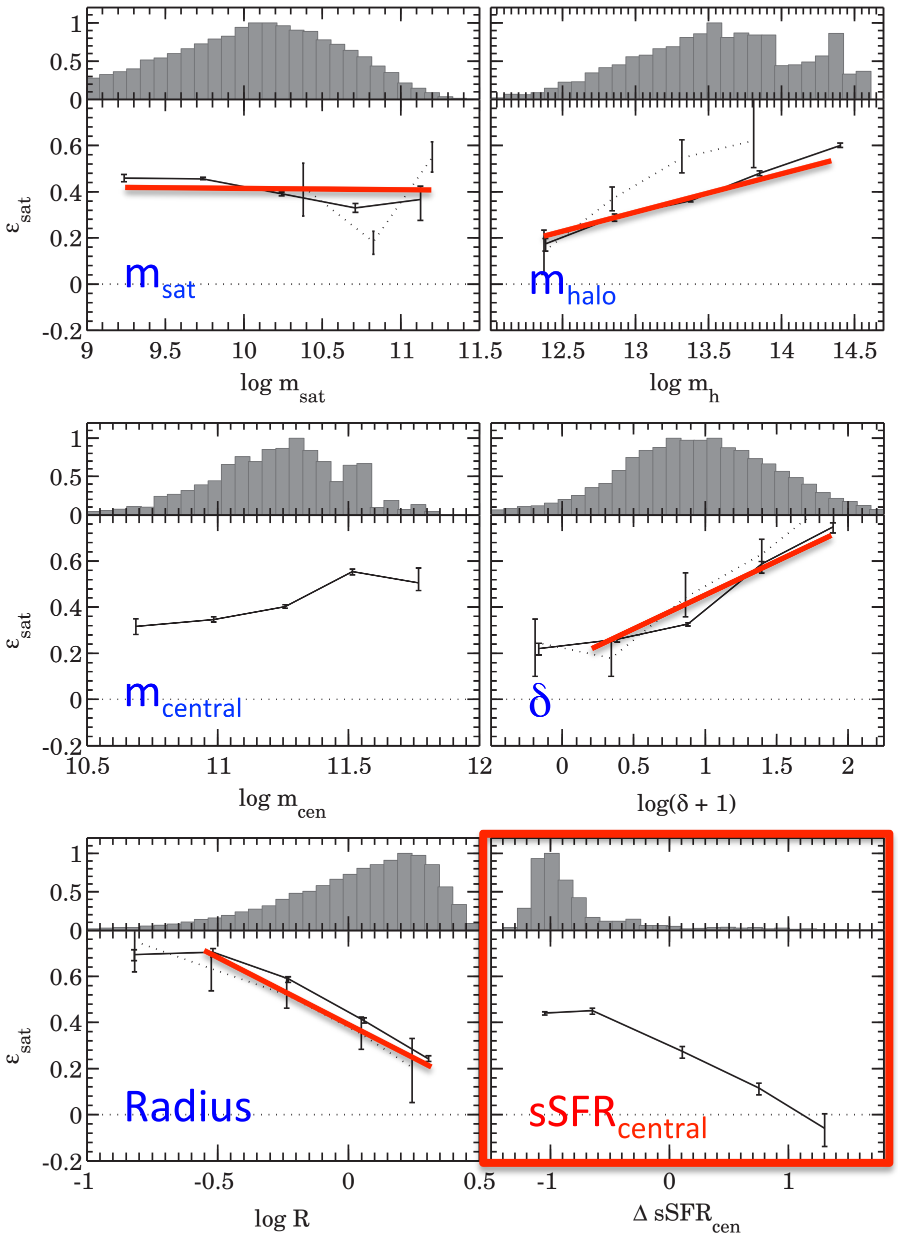

“Conformity”

Weinmann et al 2006

Variation of ε_{sat} in 2-d

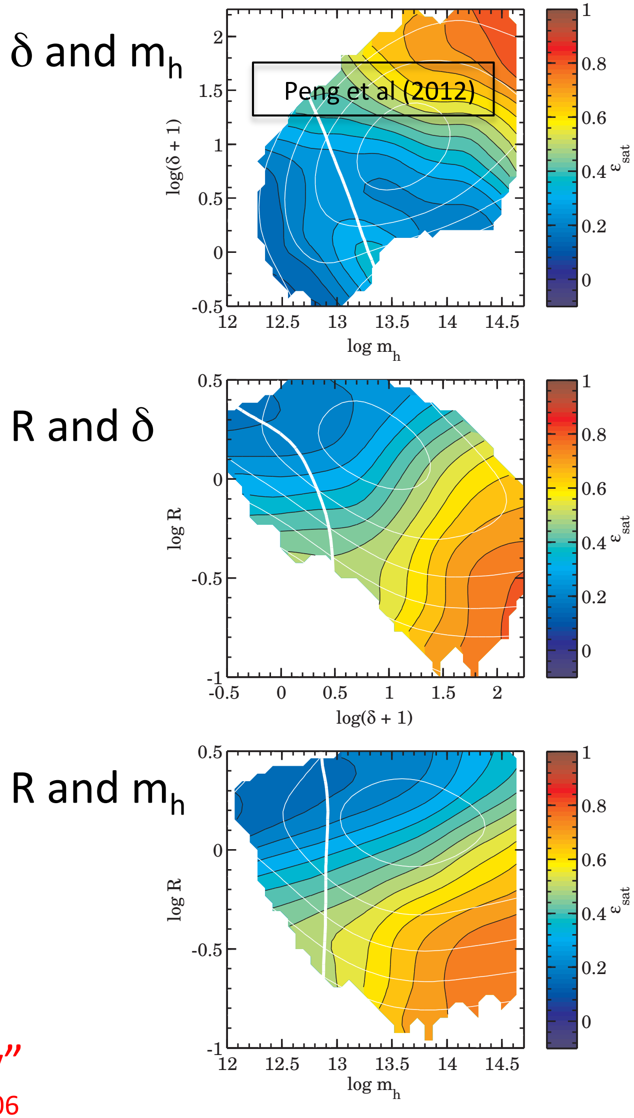

Now we can start to understand (partly) why ε_{sat} is so strikingly independent of the stellar mass of the satellite.

Distribution of most of the environment variables for satellites is <u>independent</u> of the satellite's mass (related to self-similarity of halo assembly).

But this still requires the "response" of a satellite to the quenching stimulus be independent of mass.

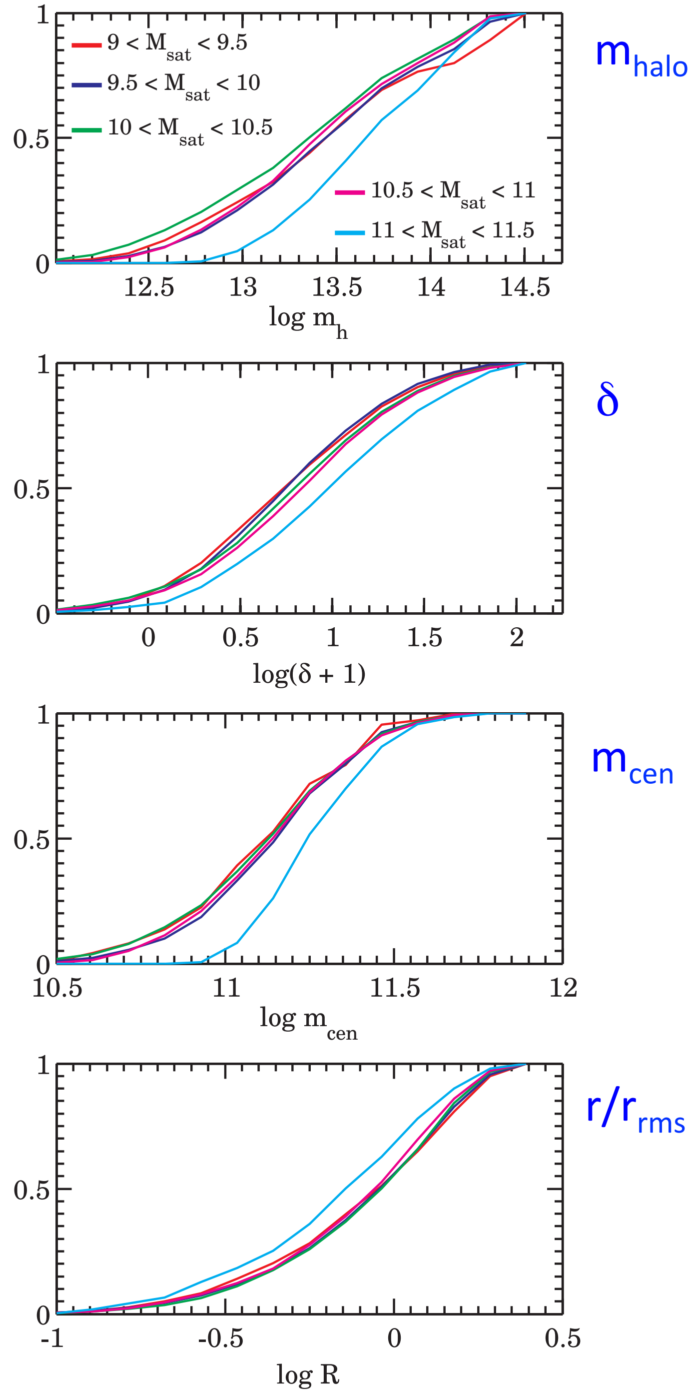

Topic 4:
Conformity: knowledge and ignorance

Sin et al (2019), MNRAS, 488, 234

Conformity introduced to describe the fact that the quenched fraction of satellites is correlated with the quenched state of their central *even when samples are matched in halo mass* Weinmann et al (2006)

Note that a correlation between quenching of satellites and centrals can arise quite trivially if the quenching of centrals and satellites both depend on some parameter which is "shared" between a given central and its satellites (e.g. halo mass)

But it is easy to see that such a correlation will disappear if the quenching of the satellites is studied at a single value of that parameter or, equivalently, if the two samples of satellites (with and without quenched centrals) are carefully matched in that parameter (as Weinmann et al did for m_{halo}), i.e. that parameter is "unhidden" or "exposed" in the analysis.

The existence of Weinmann's "conformity" in a sample tells us that there is still an additional unknown "*hidden common variable*" that is: (Knobel et al 2015)

- effecting the quenching of both satellites and centrals
- shared in some way by the central and satellite
- still hidden (in the sense of "not-matched") so that distributions are in fact different for satellites of SF and quiescent centrals
- "orthogonal" to the currently exposed variables (also measurement errors)

Note that, in this sense, "conformity" is a property of an analysis (which variables remain hidden) and NOT a physical property of the Universe. If we knew everything, conformity would disappear.

Put more generally… (first generality: forget centrals and satellites)

We could imagine some function of observable variables $g(\mathbf{v})$ that would enable us to exactly predict whether a given galaxy was quenched or not.

$$q = g(\mathbf{v}); \quad q = 0,1$$

Unfortunately we neither know *g,* nor do we have full observational knowledge of all the relevant $\mathbf{v}$, but rather only a subset $\tilde{\mathbf{v}}$ of parameters.

We can then make a prediction of the state of each galaxy using some empirically determined function *f* of these parameters, which can take any value between 0 and 1. By choosing *q* to be 0 or 1, we can make *f* simply equal to the quenched fraction f_q at that point in $\tilde{\mathbf{v}}$-space. We can then compare this empirical "prediction" with the actual observed quenched state *q* to yield, for each object a "discrepancy" Δ.

$$\Delta = q - f_q(\tilde{\mathbf{v}})$$

Δ therefore reflects our *inability* to exactly predict** the star-formation state of a galaxy. As our knowledge improves, because we have more observational information and/or a better model for how these properties affect the evolution of our galaxy, then Δ will reduce. Note that the average Δ must be zero (if *f* is set to f_q).

** *Note in context of the earlier Σ discussion: successfully "predicting" q is not the same as physically understanding the causality.*

Sin et al (2019), MNRAS, 488, 234

Put more generally…

Every galaxy will therefore have a value of Δ. We can then examine the statistical properties of Δ, specifically the variance σ_Δ^2 and the correlation function ξ(R) (for all, or any subset of galaxies):

$$\sigma_\Delta^2 = \frac{1}{N}\sum \Delta^2 = \left\langle f_q(1 - f_q)\right\rangle$$

$$\xi(R) = \frac{1}{N_p}\sum \Delta(\vec{x})\Delta(\vec{x}') \quad \text{where } |\vec{x} - \vec{x}'| = R$$

- The variance σ_Δ^2 therefore tells us the level of our ignorance from all sources,
- The correlation function ξ(*R*) tells us the contribution to our ignorance from variables that are correlated on the spatial scale *R*, i.e. conformity i.e. the physical range of environmental effects.
- The ratio γ_Δ tells us the relative importance of correlated (environmental) variables and uncorrelated (internal) variables.

$$\gamma_\Delta = \frac{\xi_\Delta\left(\lim R \to 0\right)}{\sigma_\Delta^2}$$

Sin et al (2019), MNRAS, 488, 234

A simple example…

Imagine quenching is perfectly determined from $\mathbf{v}_1$ (an internal, spatially uncorrelated variable) and $\mathbf{v}_2$ (an environmental, spatially correlated variable) by some arbitrary step-function of both $\mathbf{v}_1$ and $\mathbf{v}_2$.

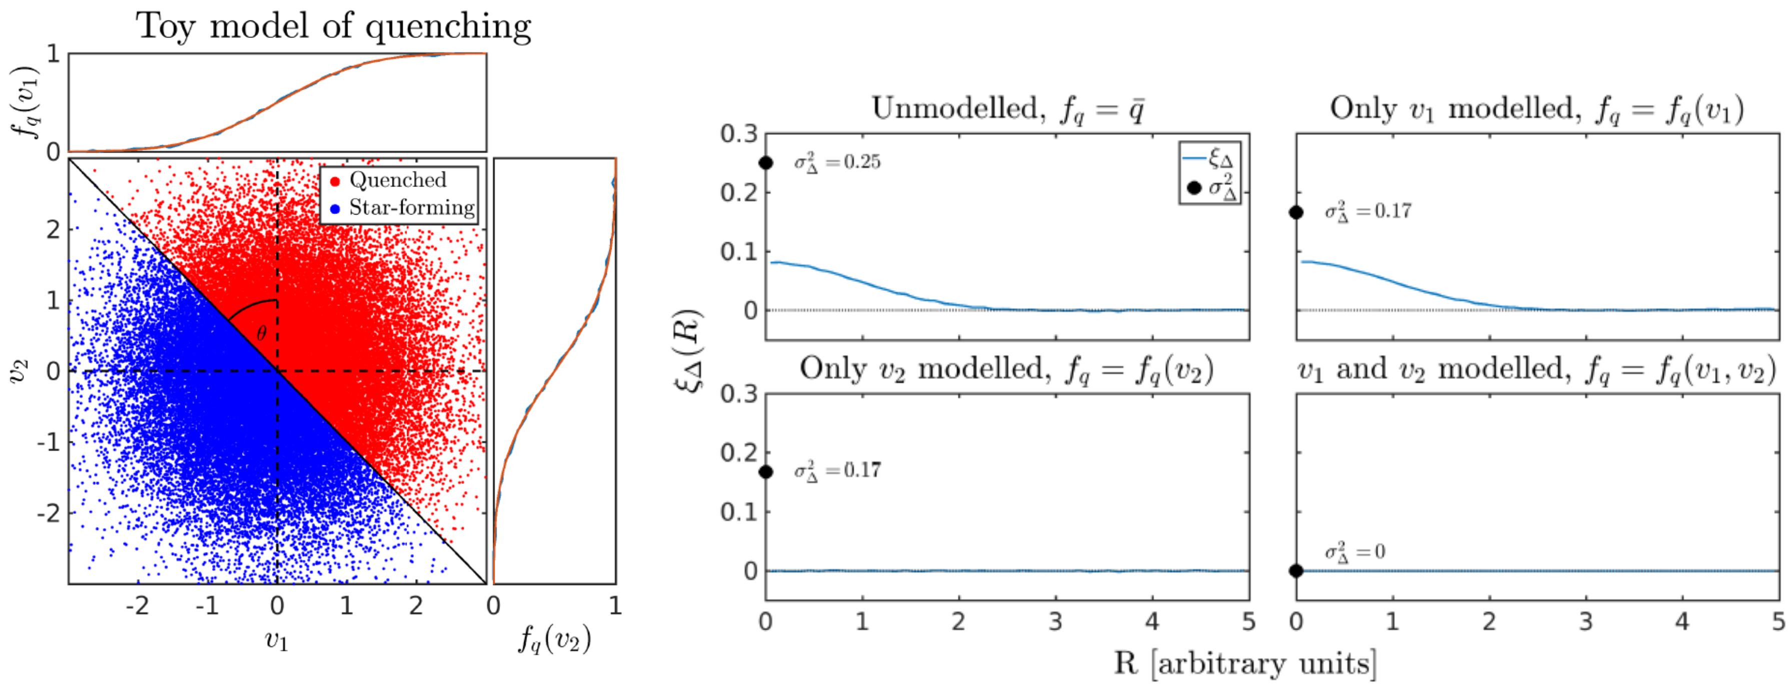

Applied to real data…

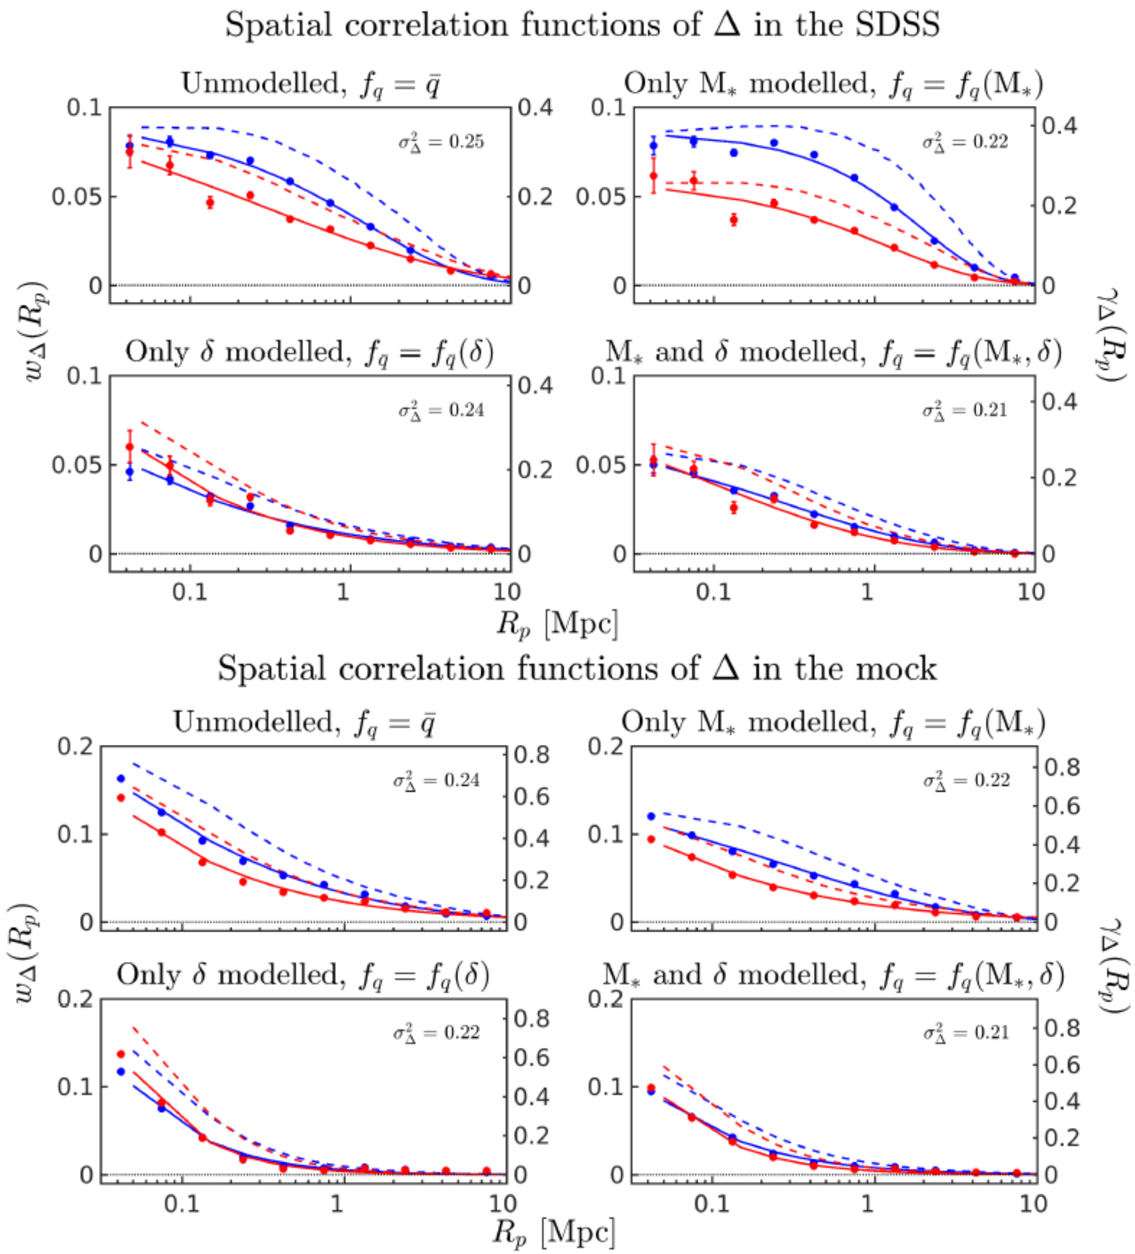

Now we can apply this to the (real) SDSS galaxy sample and to a large "mock" galaxy sample that was produced by a semi-analytic model (SAM) (for which we know the input physics).

We see strong residual conformity in both samples, at roughly comparable levels, even as both M_{star} and δ are included in the analysis.

Obvious possibility: something to do with individual haloes, going beyond $m_{\rm star}$ and simple density δ.

Sin et al (2019), MNRAS, 488, 234

Applied to real data…

So, now do the sum for the correlation function only on pairs of galaxies thought to be in different haloes. Conformity completely disappears in the mock. It had to, because there was no physics put in to the SAM on super-halo scales.

But, significant conformity *remains* in the SDSS. Is this evidence for super-halo physics? Or (we are almost certain), is it due to non-fidelity of the identification of galaxies to haloes (i.e. "group-finding").

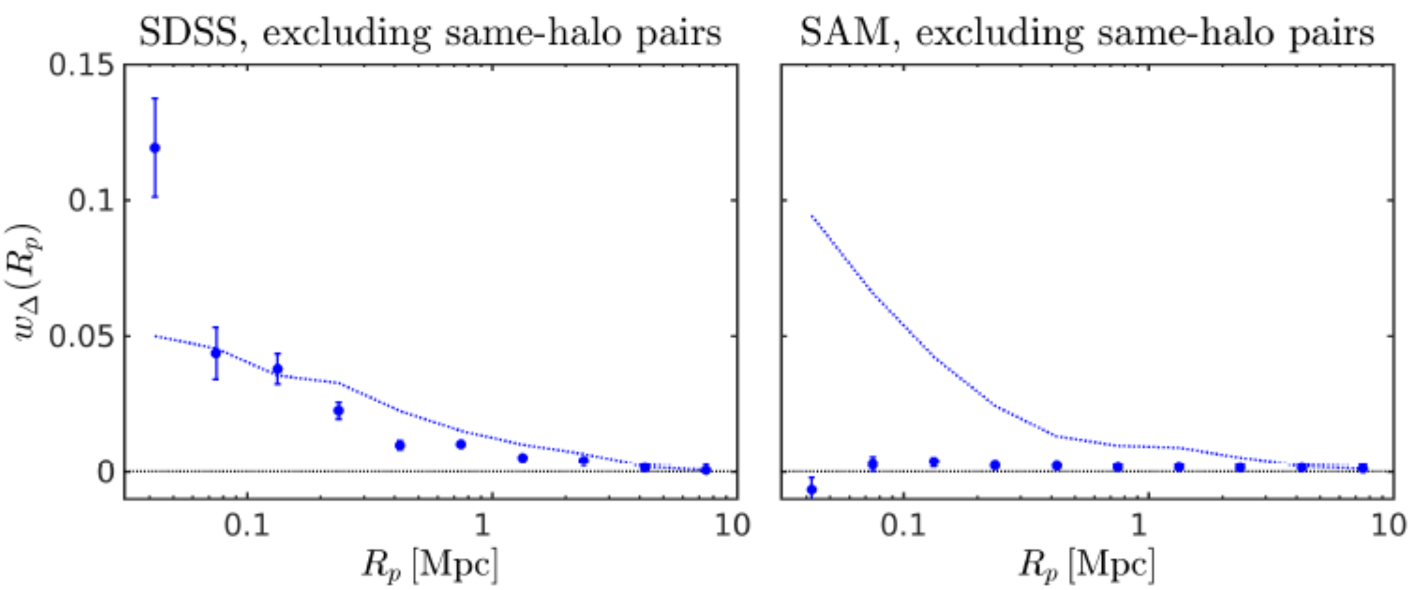

We hope this will be a good formalism to develop further….