

Annual Review of Astronomy and Astrophysics
The Circumgalactic Medium

Jason Tumlinson,^{1,2} Molly S. Peeples,^{1,2}
and Jessica K. Werk³

¹Space Telescope Science Institute, Baltimore, Maryland 21218; email: tumlinson@stsci.edu, molly@stsci.edu

²Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218

³University of Washington, Seattle, Washington 98195; email: jwerk@uw.edu

Annu. Rev. Astron. Astrophys. 2017. 55:389–432

First published as a Review in Advance on June 28, 2017

The *Annual Review of Astronomy and Astrophysics* is online at astro.annualreviews.org

<https://doi.org/10.1146/annurev-astro-091916-055240>

Copyright © 2017 by Annual Reviews.
All rights reserved

Keywords

gas, galaxies, galaxy evolution, cosmology

Abstract

The gas surrounding galaxies outside their disks or interstellar medium and inside their virial radii is known as the circumgalactic medium (CGM). In recent years this component of galaxies has assumed an important role in our understanding of galaxy evolution owing to rapid advances in observational access to this diffuse, nearly invisible material. Observations and simulations of this component of galaxies suggest that it is a multiphase medium characterized by rich dynamics and complex ionization states. The CGM is a source for a galaxy's star-forming fuel, the venue for galactic feedback and recycling, and perhaps the key regulator of the galactic gas supply. We review our evolving knowledge of the CGM with emphasis on its mass, dynamical state, and coevolution with galaxies. Observations from all redshifts and from across the electromagnetic spectrum indicate that CGM gas has a key role in galaxy evolution. We summarize the state of this field and pose unanswered questions for future research.



**ANNUAL
REVIEWS Further**

Click [here](#) to view this article's online features:

- Download figures as PPT slides
- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

Contents

1. A VERY BRIEF HISTORY	390
2. GALAXIES IN GASEOUS HALOS.....	391
2.1. The Major Problems of Galaxy Evolution	391
2.2. Our Point of View.....	395
3. HOW WE STUDY THE CGM	396
3.1. Transverse Absorption-Line Studies.....	396
3.2. Stacking Analyses.....	397
3.3. Down the Barrel.....	397
3.4. Emission-Line Maps.....	398
3.5. Hydrodynamic Simulations	398
4. THE PHYSICAL STATE OF THE CGM	399
4.1. The Complex, Multiphase CGM.....	399
4.2. From Basic Observables to Physical Properties.....	401
4.3. Line Profiles and Gas Kinematics	403
4.4. Challenges in Characterizing the Multiphase CGM	404
4.5. Gastrophysical Models.....	405
5. THE BARYONIC MASS DISTRIBUTION OF THE CGM	406
5.1. The Missing Baryons Budget.....	406
5.2. CGM Masses by Phase	407
6. METALS: NATURE’S TRACER PARTICLES	412
6.1. The Metals Census	412
6.2. Metals Observed as Gas.....	413
6.3. Metals Observed as Dust.....	415
7. INFLOWS, OUTFLOWS, AND RECYCLING	416
7.1. The Problems: Galaxy Fueling and “Missing” Metals.....	416
7.2. Empirical Signs of Fueling and Inflows	417
7.3. The Preeminence of Outflows	418
7.4. Following the Metals: The Role of Recycling	419
8. THE PARADOX OF QUENCHING.....	421
8.1. The Fate of Cold Accretion and the Problem with Recycling	421
8.2. The CGM of AGNs and Quasars	423
9. OPEN PROBLEMS, FUTURE PROSPECTS, AND FINAL THOUGHTS.....	424
9.1. Progress and Problems	424
9.2. Future Prospects for Data.....	426
9.3. Final Thoughts.....	427

1. A VERY BRIEF HISTORY

In the mid-1950s, Guido Münch observed neutral sodium (NaI) and singly ionized calcium absorption (CaII) in the spectra of hot stars at high Galactic latitudes. Before these data were published as Münch & Zirin (1961), Münch showed them to Lyman Spitzer, who interpreted the lines as evidence for diffuse, extraplanar hot gas ($T \sim 10^6$ K), which keeps the colder clouds traced by NaI and CaII in pressure confinement (Spitzer 1956). Thus was born the idea of the Galactic corona and its exploration by absorption lines in the spectra of background objects. Following Schmidt’s

1963 discovery of quasars, studies of extragalactic gas rapidly progressed with spectroscopy of the intervening absorption lines by J. Bahcall, M. Burbidge, J. Greenstein, W. Sargent, and others. Bahcall & Spitzer (1969, p. L63) then proposed that “most of the absorption lines observed in quasi-stellar sources with multiple absorption redshifts are caused by gas in extended halos of normal galaxies.” In the 1980s, subsets of the quasi-stellar object (QSO) absorption lines were associated with galaxies (Bergeron 1986, Bergeron & Boissé 1991), whereas the Lyman alpha forest emerged as their intergalactic medium (IGM) counterpart (Sargent et al. 1980). Spurred by these developments, the *Hubble Space Telescope* and the Keck Observatory made great leaps in the 1990s toward a broader characterization of the number density and column density distribution of the IGM and circumgalactic medium (CGM) back to $z \sim 3$. Pioneering studies from *Hubble*’s Key Project on QSO absorption lines demonstrated that galaxy halos give rise to strong $\text{Ly}\alpha$, CIV, and other metal lines (e.g., Lanzetta et al. 1995, Chen et al. 1998) in a gaseous medium that is richly structured in density, temperature, and ionization (**Figure 1**).

In the 2000s, large galaxy surveys such as the Sloan Digital Sky Survey (SDSS) uncovered the galactic baryon deficit, the mass metallicity relation, and quenching problems (Section 2). Meanwhile, theorists implemented new physical prescriptions for gas accretion and feedback with new numerical methods and faster computers. It soon became impossible even to address these big mysteries of galaxies without appealing to gas flows between the interstellar medium (ISM), the IGM, and by implication, the CGM. Yet most such models of gas flows were, and are still, tested against observations of starlight—the same observations that first posed the problems. By the mid-2000s, models and observations of gas flows in and out of galaxies had reached the point that the former were in urgent need of direct observations of the gas flows themselves. CGM studies leaped forward in the late 2000s with the installation of *Hubble*’s *Cosmic Origins Spectrograph*, which was designed for reaching diffuse gas with $30\times$ the sensitivity of its predecessors, and with new techniques for stacking and combining X-ray and optical spectra. This is the context in which we write our review. We aim to survey recent progress in observing and modeling the gas flows that drive galaxy evolution and, thus, to tell the story of galaxy evolution writ large from the perspective of the CGM.

For additional perspective on the issues raised here from a more Galactic point of view, we recommend the recent Annual Review article on halo gas by Putman et al. (2012). For an up-to-date survey of accretion, see the forthcoming volume *Gas Accretion onto Galaxies* (Fox & Davé 2017).

CGM: circumgalactic medium

IGM: intergalactic medium

ISM: interstellar medium

SDSS: Sloan Digital Sky Survey

2. GALAXIES IN GASEOUS HALOS

2.1. The Major Problems of Galaxy Evolution

We motivate and organize our review with four major galaxy evolution problems in which the CGM is implicated (**Figure 2**). We consider these problems in connection with the simulated CGM illustrated in **Figure 3**. Why do dark matter halos of different masses give rise to galaxies with drastically different star-formation and chemical histories (Sections 2.1.1 and 2.1.2)? Why does such a small fraction of cosmic baryons and metals reside in the galaxies (Sections 2.1.3 and 2.1.4)? The prevailing answers to these questions all feature the regulation of gas flows into and out of galaxies—which necessarily pass into and through the CGM. We initially pose these problems at low redshift, but they all have high- z counterparts, and their solutions require understanding the CGM and the flows that feed it at all cosmic epochs.

2.1.1. How do galaxies sustain their star formation? Star-forming galaxies pose a conundrum: Their ISM gas can last for only a small fraction of the time they have been forming stars (**Figure 2a**), implying an external supply of gas that keeps the ISM in a quasi-equilibrium state. The depletion




Figure 1

A diagram of the CGM. The galaxy's red central bulge and blue gaseous disk are fed by filamentary accretion from the IGM (blue). Outflows emerge from the disk in pink and orange, whereas gas that was previously ejected is recycling. The diffuse gas halo in varying tones of purple includes gas that is likely contributed by all these sources and mixed together over time. Refer to **Supplemental Figure 1** for an alternate version of the figure, which illustrates the different observing techniques we discuss in Section 3. Abbreviations: CGM, circumgalactic medium; IGM, intergalactic medium.

Supplemental Material

time, $\tau_{\text{dep}} \sim M_{\text{gas}}/\dot{M}_{\text{sfr}}$ changes by only $\sim 2\times$ over the factor of 30 between sub- L^* and super- L^* galaxies. More generally, sub- L^* galaxies generally have extended bursty star-formation histories, as opposed to the more continuous star formation found in more massive galaxies, suggesting differences in how and when these galaxies acquire their star-forming fuel. As this fuel is from the CGM, we must explain how sub- L^* and L^* galaxies fuel star formation for longer than their τ_{dep} .

2.1.2. What quenches galaxies, and what keeps them that way? How galaxies become and remain passive is one of the largest unsolved problems in galaxy evolution (**Figure 2b**). Proposed solutions to this problem involve controlling the gas supply, either by shutting off IGM accretion or keeping the CGM hot enough that it cannot cool and enter the ISM. Low-mass galaxies tend to continue forming stars unless they are a satellite of or near a more massive galaxy (Geha et al.





Figure 2

Four important problems in galaxy evolution viewed with respect to M_* . (a) The gas depletion timescale $\tau_{\text{dep}} \sim M_{\text{gas}}/\dot{M}_{\text{sfr}}$ for star-forming galaxies at $z \sim 0$, with M_{gas} from Peeples et al. (2014) and \dot{M}_{sfr} from Whitaker et al. (2012); the shading denotes ± 0.15 dex scatter in \dot{M}_{sfr} . (b) The galaxy bimodality in terms of M_* and specific SFR (Schiminovich et al. 2010). (c) The galactic baryon fraction, $M_*/[(\Omega_b/\Omega_m)M_{\text{halo}}]$ from Behroozi et al. (2010), with stars in red and interstellar gas in blue. (d) The “retained metals fraction,” metals for several galactic components relative to all the metals a galaxy has produced, with stars in red, interstellar gas in blue, and interstellar dust in orange. Adapted from Peeples et al. (2014) with permission. Vertical bars mark the properties of sub- L^* , L^* , and super- L^* galaxies at $\log M_*/M_\odot = 9.5$ (blue), 10.5 (green), and 11.0 (red), respectively. Abbreviation: SFR, star-formation rate.



2012). This finding suggests that the central galaxy’s gaseous halo strips the satellite with ram pressure or “starves” the satellite of fresh fuel. These ideas have specific testable implications for the physical state of the CGM.

2.1.3. Why do galaxies lack their fair share of baryons? The Λ CDM model predicts that baryons follow gravitationally dominant dark matter into halos, where the gas dissipates energy as radiation and cools into the center of the halo. Observed galaxies, however, harbor only a small share of the halo’s expected baryons in their stars and ISM, with $M_b \ll (\Omega_b/\Omega_m)M_h$ (**Figure 2c**). Even at their most “efficient,” L^* galaxies have converted only $\sim 20\%$ of their halos’ baryons into stars (**Figure 2c**), with values of only about 5–10% in sub- L^* and super- L^* galaxies (Behroozi et al. 2010, McGaugh et al. 2010). There are three basic possibilities: The baryons are in the halo but not yet detected, such as hot or diffuse gas; the baryons have been accreted and then ejected from the halo altogether; or the baryons have been prevented from accreting into the halo in the first place. Although reality probably combines aspects of all three, in any combination they strongly suggest that the CGM is an excellent place to look for missing halo baryons in cold or hot gas or to look for direct evidence of past ejection.

2.1.4. Where are the metals? Although baryons come from outside the halo, metals are sourced locally by stars and the deaths of stars. Star-forming galaxies over ~ 3 decades in stellar mass retain a surprisingly flat ~ 20 – 25% of the metals they have ever produced (Peeples et al. 2014) in their stars, ISM gas, and dust. Metals have clearly been lost to outflows (Tremonti et al. 2004), but how these outflows scale with galaxy mass is unclear. Models that already struggle to reproduce the observed steep mass–metallicity relation (Somerville & Davé 2015) fail to retain the low, flat fraction of metals produced (e.g., Zahid et al. 2012, Muratov et al. 2015, Oppenheimer et al. 2016). Although “missing baryons” concern accretion and feedback through the outer boundary of the CGM, metals address the disk/halo interface: Do they leave the halo altogether or recycle back into the galaxy’s ISM on long timescales as a “halo fountain” (Oppenheimer & Davé 2008)? On what timescales are ejected metals recycled? How metal-enriched is outflowing material relative to the ambient ISM; i.e., what are the entrainment fractions and metal-loading factors? How does dust survive the journey out of galaxies, and what chemical clues does it hide? As we show below, following the metals as “Nature’s tracer particles” is a fruitful and revealing route to understanding the CGM.

2.2. Our Point of View

How galaxies acquire, eject, and recycle their gas are core issues in galaxy evolution, on par with how they evolve in their shapes and how star formation works. To a large extent these gas flows are galaxy evolution. The CGM is a main venue for these flows: It is potentially the galactic fuel tank, waste dump, and recycling center all at the same time. This review approaches the growing body of empirical results and theoretical insights from the direction of these four major questions. For example, rather than asking “What are the MgII absorbers?,” we ask, “What do the MgII absorbers

Λ CDM:
cold-dark-matter
cosmology with a
cosmological constant


 **Supplemental Material**

Figure 3

These simulated views (from EAGLE; adapted from Schaye et al. 2015 and Oppenheimer et al. 2016) of the CGM are more sophisticated but possibly just as uncertain as **Figure 1**. The three columns render a single galaxy with $M_* = 2.5 \times 10^{10} M_\odot$ at $z = 0$ in density (*left*), temperature (*middle*), and metallicity (*right*). The galaxy is shown at redshifts $z = 3, 2, 1$, and 0 from top to bottom. The dotted white circle encloses the virial radius at each epoch. Here, each panel is 1 Mpc physical across; refer to **Supplemental Figure 2** for alternate versions of the figure scaled to 1 Mpc comoving and to within the virial radius.

Phenomenology:
emergent properties
and scaling relations

Physics: underlying
physical properties and
processes

X-ray: $\lambda \lesssim 30 \text{ \AA}$

**Extreme-ultraviolet
(EUV):**
 $400 \lesssim \lambda \lesssim 900 \text{ \AA}$

**Far-ultraviolet
(FUV):**
 $900 \lesssim \lambda \lesssim 2000 \text{ \AA}$

**Near-ultraviolet
(NUV):**
 $2000 \lesssim \lambda \lesssim 3400 \text{ \AA}$

tell us about the mass and kinematics of galactic outflows?” We thus favor physical insights and synthesis of discoveries over detailed discussions of methods, compilations of data, or exhaustive cataloging of the literature. We hope that this approach improves understanding between those who study gas and galaxies (still disparate groups) and those who more effectively highlight open issues to be pursued in the future.

For the purposes of our discussion, we define the CGM to be bounded at the outside by the virial radius R_{vir} of a galaxy’s dark matter halo and on the inside by the disk or ISM. Neither boundary is well defined, and precisely defining when gas passes through one of these boundaries can be either a valuable research contribution or a fruitless semantic exercise depending on circumstances. We focus on the phenomenology and physics of gas that fills out halos without too much attention to these exact definitions.

3. HOW WE STUDY THE CGM

We now survey the wide range of observational techniques used to probe the CGM. We include diagnostic features from across the electromagnetic spectrum from the X-ray through the extreme-ultraviolet (EUV), far-ultraviolet (FUV), and near-ultraviolet (NUV), and into the optical.

3.1. Transverse Absorption-Line Studies

Viewing the CGM in absorption against a bright background source like a quasar offers three major advantages over other methods: (a) sensitivity to extremely low column density, $N \simeq 10^{12} \text{ cm}^{-2}$; (b) access to a wide range of densities, unlike emission-line measures that scale as density squared; and (c) invariance of detection limits to redshift and the luminosity of the host galaxy. These advantages come at a cost, however: Absorption provides only projected, pencil-beam measures of gas surface density, usually limited to one sightline per galaxy by the rarity of background quasars. Within the local Universe (a few megaparsecs) it is possible to use multiple sightlines (e.g., Lehner et al. 2015, Bowen et al. 2016) and, at higher redshift, multiply lensed images from background quasars (e.g., Rauch & Haehnelt 2011, Rubin et al. 2015) to constrain the sizes of absorbers. In general, however, CGM maps made from absorption-line measurements are a statistical sampling of gas aggregated from many galaxies. With massive optical spectroscopic surveys, samples have grown to hundreds or thousands in low ions like MgII and CaII (e.g., Zhu & Ménard 2013b). Quasar/galaxy pairings have now been extended out to $z \sim 4$ and beyond (Matejek & Simcoe 2012, Turner et al. 2014).

There are three basic ways of building absorber samples. First, blind surveys select background quasars based on brightness and/or redshift and so are optimal for samples that are unbiased with respect to foreground structure. Ground-based redshift surveys around previously observed quasar sightlines are now a time-honored method for constructing samples of quasar/galaxy pairs (e.g., Chen et al. 1998, Stocke et al. 2006, Rudie et al. 2012). The second, targeted approach chooses background sources because they probe particular foreground structures, such as L^* galaxies (Tumlinson et al. 2013), sub- L^* galaxies (Bordoloi et al. 2014b), galaxies with known ISM content (Borthakur et al. 2015), or groups and filaments (Wakker et al. 2015, Tejos et al. 2016), by cross-matching the observable quasar with catalogs of these structures. Finally, maps of absorption in the Milky Way (MW)’s CGM use essentially any quasar (or UV-bright halo stars), sometimes chosen to pass through known halo gas structures and sometimes not. Though most absorption-line work has been in the UV and optical, the *Chandra X-ray Observatory* and *X-ray Multi-Mirror Mission-Newton* (XMM-Newton) have been used to search for X-ray gas in individual absorbers, constraining the extent of CGM and IGM hot gas (Nicastro et al. 2005).

It is useful to distinguish between HI column density regimes that must be, or can be, treated differently in analysis. Lines up to $\log N \simeq 15$ can usually be analyzed with equivalent widths or Voigt profile fitting. The value $\log N \simeq 15$ is high for the Ly α forest but low for the CGM [there are, of course, a few exceptions (Tumlinson et al. 2013, Johnson et al. 2014) in which HI is not seen at <100 kpc even to low limits]. At $\log N \simeq 16$ saturation becomes a major factor, and robust column densities (as opposed to lower limits) must come from profile fitting or from the higher Lyman series lines if the system is redshifted enough. Systems with $\log N \simeq 16$ are partial or complete Lyman limit systems (LLSs). If the Lyman limit is covered ($z > 0.24$ for *Hubble*), the flux decrement at $\lambda = 912(1+z)$ Å allows a precise measurement of $\log N_{\text{HI}}$ and improved ionization and metallicity diagnostics. Above $\log N_{\text{HI}} \simeq 18$ (where N_{HI} is the HI column density in cm^{-2}), the Lyman limit is totally opaque, the highest Lyman series lines are saturated, and genuine column densities must come from fitting the Ly α profile for LLSs and damped Lyman- α systems.

Lyman limit system (LLS):

$N_{\text{HI}} > 10^{16.2} \text{ cm}^{-2}$,
the “dense” CGM

3.2. Stacking Analyses

Massive spectroscopic surveys have enabled another novel method for examining halo gas. “Stacking” of hundreds or thousands of spectra is a powerful way to extract faint signals from absorption-line data sets. This technique requires catalogs of redshifts, for either foreground galaxies or absorbers, so that the spectra of background objects can be shifted to their rest frames and continuum-normalized and then co-added together. The co-addition beats down statistical noise, enabling measurements of weak absorption at the cost of averaging over individual absorber profiles. When the catalogs of foreground galaxies include properties such as mass, radius, SFR, color, environment, or orientation, the stacks can be performed with subsets of the data to examine the variation of mean profiles with these properties (York et al. 2006, Bordoloi et al. 2011, Zhu & Ménard 2013a). Stacking experiments that correlate the reddening of quasars with foreground galaxy halos in the SDSS survey have revealed large quantities of dust in the CGM of galaxies (Ménard et al. 2010, Peek et al. 2015). Stacking techniques can also exploit more numerous, but fainter, sources; for example, Steidel et al. (2010) characterized the CGM of $z \sim 3$ galaxies by stacking the spectra of background galaxies. Stacking can detect weak signals in the mean properties of gas absorbers but at the cost of averaging out kinematic and ionization structure that may contain significant physical meaning.

3.3. Down the Barrel

“Down-the-barrel” spectroscopy uses a galaxy’s own starlight as a background source for detecting absorption. This method has been a fruitful one for studying galactic inflows and outflows from spectroscopy of star-forming galaxies. This method is commonly used to study inflows and outflows in galaxies out to $z \sim 2\text{--}3$ in redshifted FUV, NUV, and optical lines (Martin 2005; Steidel et al. 2010; Bordoloi et al. 2011; Kornei et al. 2012; Rubin et al. 2012, 2014; Heckman et al. 2015; Henry et al. 2015). Down-the-barrel measurements are critical pieces of the CGM puzzle because they directly trace current outflows at galactocentric radii that are inefficiently covered by background sources (because of the R^2 scaling of foreground cross-section). Although down-the-barrel spectra are key for tracing the accretion and outflows that dominate CGM kinematics, they have the key limitation that the galactocentric radius of any detected absorption is unconstrained—it could be anywhere along the line of sight—complicating mass and covering fraction estimates inferred from these spectra.

3.4. Emission-Line Maps

Emission-line observations search for photons emitted directly from CGM gas. As the emission measure scales as n^2 , and the CGM has $n_{\text{H}} \sim 10^{-2}$ or less, finding these photons is a stiff challenge. The MW halo has been extensively mapped for high-velocity clouds (HVCs) and other halo structures using radio emission at 21 cm. This technique has been applied to external galaxies (Putman et al. 2012), but detections are limited to within $\sim 10\text{--}20$ kpc of the targeted galaxies. The soft X-ray band is optimal for gas at $\gtrsim 1$ MK. The extremely low surface brightness of the gas makes these observations challenging and expensive, but a few individual halos have been detected and their hot gas budgets measured by *Chandra* and/or *Röntgensatellit* (e.g., Humphrey et al. 2011, Anderson et al. 2016). Stacking-of-individual-galaxies techniques has also yielded mass density profiles for hot gas around nearby galaxies (Anderson et al. 2013). When combined with halo size, density, and metallicity constraints from soft X-ray absorption-line techniques, these maps have aided in the assessment of the total mass and baryon fraction of the hot CGM.

Emission-line maps are also possible at UV/optical wavelengths, though they are no less challenging than in the X-ray band. Recent reports claim a detection of an extended OVI halo ($R \sim 20$ kpc) around a low-redshift starburst galaxy (Hayes et al. 2016). Extended Ly α emission has been seen out to ~ 100 kpc away from $z \sim 2.5$ galaxies and QSOs (Cantalupo et al. 2014, Prescott et al. 2015). In another case, an extended filamentary structure connected to a galactic disk was detected using diffuse emission in the optical band (Martin et al. 2015). Emission maps can constrain the density profile, morphology, and physical extent of the gas more directly than aggregated pencil-beam sightlines (Corlies & Schiminovich 2016). For X-ray emission from fully ionized gas, masses can be inferred more directly, avoiding the uncertain ionization corrections that plague absorption-line measurements (Section 4); indeed, the CGM’s more massive cousin, galaxy clusters’ intracluster medium, has been studied in detail via X-ray emission for decades (Vikhlinin et al. 2006). On the downside, emission-line maps are still challenging technically; the surface brightnesses are extremely small compared to sky and detector backgrounds, and surface brightness dimming has a steep increase with redshift. In a recent study using stacks of fiber spectra from SDSS, Zhang et al. (2016) achieved detections of H α at 50–100 kpc around low-redshift galaxies, demonstrating that very sensitive limits can be reached on galaxies in the aggregate. These observations remain challenging, but as “taking a picture” of an astrophysical object remains the ideal, efforts to improve instrument technology and enable emission-line mapping to reach samples of hundreds of galaxies across cosmic time is an important goal.

3.5. Hydrodynamic Simulations

Physical models and simulations are essential tools for understanding the CGM. In contrast to observations, they provide for controlled environments in which physical properties, histories, and futures of gas are all known and can be manipulated to tease insights out of the otherwise unmanageable complexity of a multiphase gaseous medium. As reviewed by Somerville et al. (2015), there are many schemes for simulating the development of the cosmic web and galaxies under the influence of dark matter, gravity, and hydrodynamics. The major methods at present are smoothed particle hydrodynamics, such as Gadget (Ford et al. 2013, Oppenheimer et al. 2016), Gasoline (Christensen et al. 2016, Gutcke et al. 2017), and GIZMO (Muratov et al. 2017); adaptive mesh refinement, such as Enzo (Hummels et al. 2013, Corlies & Schiminovich 2016); and moving mesh (Arepo and the Illustris simulation; Suresh et al. 2017). Large-scale cosmological simulations in megaparsec-scale boxes can simulate hundreds of galaxies in their proper Λ CDM contexts (e.g., Oppenheimer & Davé 2006, Ford et al. 2014, Vogelsberger et al. 2014). At the opposite end of the scale, very high-resolution simulations focused on the interaction between dense clouds and diffuse

halos (e.g., Heitsch & Putman 2009, Armillotta et al. 2016) that can reach scales at much less than a parsec. Spanning these two regimes are the so-called zoom simulations, which resolve enough of the large-scale structure to accurately trace a single galaxy or a subset of galaxies selected out of larger boxes (**Figure 3**; Schaye et al. 2015). Even zooms must make assumptions about physics that they do not resolve, using “sub-grid” prescriptions to stand in for such complex phenomena as star formation, metal mixing and transport, supernova and AGN feedback, and others. Sub-grid models are parameterized and tuned to yield specific metrics—like the stellar mass function at $z = 0$ —and then the properties that emerge—such as SFRs, morphology, quenching, and the CGM—are analyzed and compared with data to constrain the physical prescriptions that went in. We use simulations from a broad range of techniques and groups to look for insights into how the CGM participates in galaxy evolution and to help interpret data.

4. THE PHYSICAL STATE OF THE CGM

We now turn to the density profile, phase structure, and kinematics of the CGM. We first present the data that show the various ionization states and velocity distributions of the CGM absorption (Section 4.1). Next, we describe how the absorption-line measurements may be translated into physical parameters such as density, temperature, and size (in Section 4.2). We then draw lessons from kinematics (Section 4.3) before considering the physical complexities and challenges inherent in the interpretation of these data (Sections 4.4 and 4.5).

4.1. The Complex, Multiphase CGM

As a matter of empirical inference, the CGM is multiphase in its ionization structure and complex in its dynamics. The ionization structure shown in **Figure 4** compiles measurements for six diagnostic ions as a function of impact parameter (a proxy for radius). These data indicate a wide range of density and ionization conditions up to a few 10^5 K with very little interpretation required. Observationally, multiphase means many of these metal ions spanning an order of magnitude in ionization potential energy are commonly found within the same absorber system occupying a galaxy’s halo. An open question in the physics of circumgalactic gas is what this observed multiphase ionization structure reveals about the small-scale multiphase density, temperature, and metallicity structure of the CGM.

Over the past 20 years, the practice of using such empirical inputs in analytic arguments to infer the physical state and structure of the diffuse plasma has matured greatly (Mo & Miralda-Escudé 1996, Maller & Bullock 2004). To produce an extended, multiphase CGM, authors have proposed several scenarios, which we categorize as follows: (a) massive inward cooling flows driven by local thermal instabilities (e.g., McCourt et al. 2012); (b) boundary layers between moving cool clouds in a hot atmosphere (e.g., Begelman & Fabian 1990); and (c) the continual shocking and mixing of diffuse halo gas by galactic outflows (e.g., Fielding et al. 2017, Thompson et al. 2016). We discuss the applicability of some of these analytic models in Sections 4.4 and 4.5.

Direct evidence for a hot component ($\log T \gtrsim 6$) in the multiphase CGM comes from both diffuse soft X-ray emission (Anderson & Bregman 2010, Anderson et al. 2013) and absorption along QSO sightlines (Williams et al. 2005, Gupta et al. 2012) for the MW and external galaxies. Indirect evidence for a hot phase comes from highly ionized metals that correlate with the low-ionization HVCs (Sembach et al. 2003, Fox et al. 2006, Lehner et al. 2009, Wakker et al. 2012), suggesting boundary layers between a hot medium and the colder HVCs. MW HVCs also show head–tail morphologies indicative of cool clouds moving through a hot medium (e.g., Brüns et al. 2000). Finally, the multiphase CGM is clearly manifested in hydrodynamic simulations, which

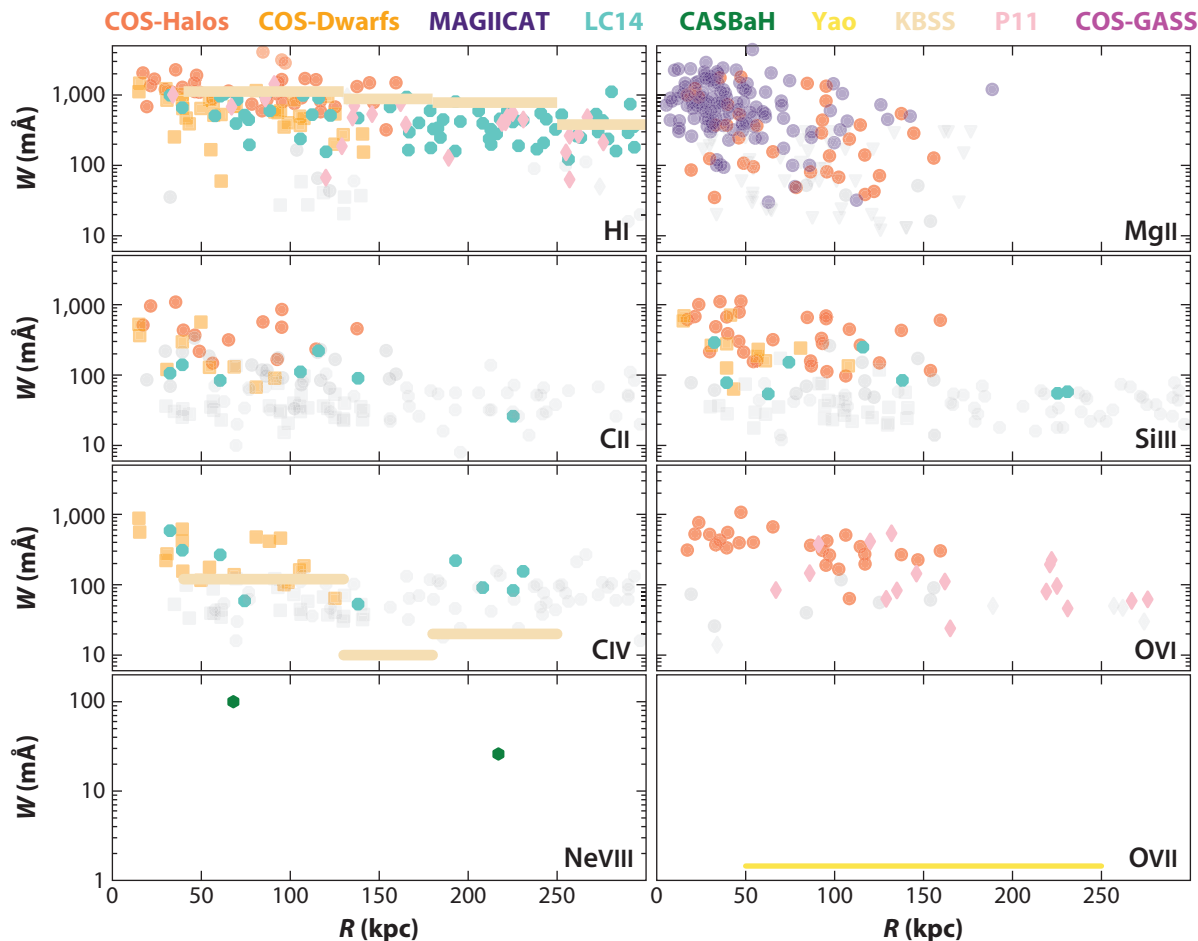


Figure 4

A range of ion equivalent-width (rest-frame) measurements for a compilation of published surveys; refer to **Supplemental Figure 3** for alternate versions showing column densities, where available. We progress from H I through seven metallic ions of increasing ionization potential. The surveys are COS-Halos (Tumlinson et al. 2013, Werk et al. 2013), COS-Dwarfs (Bordoloi et al. 2014b), COS-GASS (Borthakur et al. 2015), MAGIICAT (Nielsen et al. 2013, Liang & Chen 2014), the Keck Baryonic Structure Survey (Rudie et al. 2012, Turner et al. 2015), CASBaH (Tripp et al. 2011, Prochaska et al. 2011b), and the X-ray study of Yao et al. (2012) that imposes a stacked upper limit on O VII. Gray points indicate upper limits.

Supplemental Material

exhibit a mixture of cool (10^4 K) and warm-hot ($10^{5.5}$ – 10^6 K) gas within a galaxy virial radius with a density profile that drops with increased distance from the central host galaxy (e.g., Stinson et al. 2012, Ford et al. 2013, Shen et al. 2013, Suresh et al. 2017; **Figure 3**). For practical purposes we can regard the outer boundary of the CGM to correspond to R_{vir} , but there is no empirical reason to believe that any special behavior occurs at that radius; current observations favor trends in column densities that scale with R_{vir} but do not change in form at that arbitrary boundary.

Evidence for kinematic complexity is revealed as the detected ion species breaking into different components with distinct velocities and line widths. Shown in **Figure 5**, the various metal ions show significant but varied correspondence in their component structure. The combination of both aligned and misaligned components between ionization states may reflect clouds or streams with density structure or a population of clouds with different ionization states projected together




Figure 5

A selection of absorption-line data and Voigt profile fits from the COS-Halos survey (Werk et al. 2016), showing a range of metal ions and HI on a common velocity scale with the galaxy at $v = 0 \text{ km s}^{-1}$ on the x -axis.

along the line of sight to the same range of observed velocities. Cloud sizes are difficult to constrain in a model-independent way, but multiply lensed images from background quasars (Rauch et al. 2001, Rauch & Haehnelt 2011) prefer 1–10-kpc scales. Fitting Voigt profiles to multicomponent absorption yields column density N , a Doppler b parameter, and velocity offset v for each component from the galaxy systemic redshift, as well as the total kinematic spread of gas in a halo (but this fitting is subject to issues caused by finite instrumental resolution). Generally, the kinematic breadth of an absorber system is thought to reflect the influence of the galaxy’s gravitational potential, bulk flows, and turbulence in the CGM.

4.2. From Basic Observables to Physical Properties

We must characterize the ionization states, chemical composition, and density to properly describe the symbiotic relationship between the gas and stars in the central galaxy disk and the CGM. If it were feasible to obtain precise measurements for every ion of every abundant element, in all velocity components, then the gas flows, metallicity, and baryon budget of the multiphase CGM would be well constrained. However, atomic physics dictates that only a subset of the ionization states of each element lie at accessible wavelengths. Taking oxygen as an example, OI and OVI place strong lines in the FUV wavelengths, whereas OII–OV lines appear in the EUV wavelengths (400–900 Å). OVII and OVIII, arising in hot gas, have strong transitions in the soft X-ray ($\sim 20 \text{ Å}$) bands. Although it is therefore possible in principle to detect (or limit) every stage of oxygen, this potential has yet to be realized.

Figure 6 shows the basic schema for constraining CGM gas properties with these multiphase ions. The gray-scale phase diagram renders the properties of all $< R_{\text{vir}}$ gas from a MW-mass EAGLE zoom simulation (Oppenheimer et al. 2016). Accessible ions at each temperature and density are marked with colored squares and dashed lines. This plot is intended to be a useful guide for finding the most likely tracers of a given CGM gas phase. It cannot be used to extract precise temperatures and densities for any given ion because the metal ion positions on this phase diagram are model dependent. The inset shows the most common strong lines from these species plotted as observed wavelength versus redshift; the rest-frame wavelength is where each intercepts $z = 0$. Practically, FUV lines are available at $z < 1$ with *Hubble* and $z > 2$ from the ground, the EUV lines can be reached at $z \gtrsim 0.5$ –1 with *Hubble* ($\lambda_{\text{obs}} \gtrsim 1100 \text{ Å}$), and the X-ray lines can currently only be detected toward the small number of bright QSOs and blazars with reach of the sensitivity of *Chandra* and the XMM-Newton. As a result, most CGM measurements rely on heterogeneous ion sets—several low ions from C, N, Si, and Mg, a few intermediate ions from

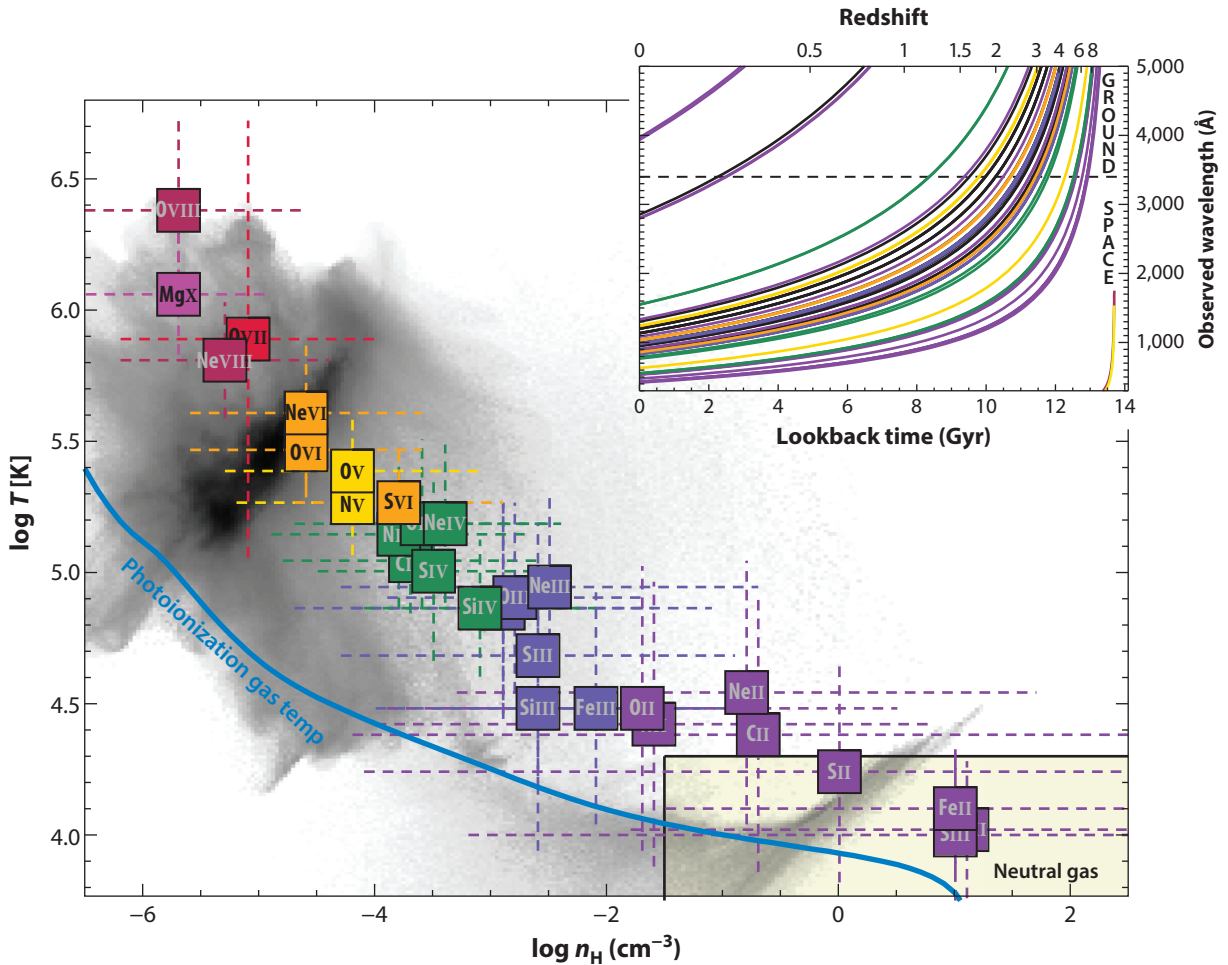


Figure 6

Metal absorption lines (ions) of the CGM from neutral to O VIII having $19 < \lambda_{\text{rest}} < 6000 \text{ \AA}$ shown on a phase (T – n_{H}) diagram within R_{vir} of the $z = 0$ EAGLE simulation shown in **Figure 2**. The points are colored according to ionization state, ranging from neutral (I, shaded yellow) to singly ionized (II, purple) to highly ionized (X, magenta). The position of each point is set on each axis where its ionization fraction peaks in CIE (temperature axis) and a standard PIE model (density axis) (Gnat & Sternberg 2007, Oppenheimer & Schaye 2013b); the range bars show the T and n range over which each species has an ionization fraction over half its maximum value (i.e., the FWHM). Complete line lists are available in Morton (2003). Refer to **Supplemental Figure 4** for alternate versions broken down by element. Abbreviations: CGM, circumgalactic medium; CIE, collisional ionization equilibrium; FWHM, full width at half maximum; PIE, photoionization equilibrium.

Supplemental Material

C and Si, and a high ion or two from Ne and O. Therefore, the gas density and temperature can only be understood in the context of a model for its ionization state (and abundance patterns).

Many assumptions are necessary to make progress toward physical models of the CGM. The two most generic classes of models are photoionization equilibrium (PIE) and collisional ionization equilibrium (CIE). Generally, low and intermediate ions can be accommodated within PIE models, whereas high ions require CIE models. Species at intermediate ionization potentials, such as C IV and O VI, sometimes show a preference for one or the other or have contributions from both.

These two classes of model are not mutually exclusive: A gas that is collisionally ionized may have the ion ratios further affected by incident radiation, and there are numerous possible departures from equilibrium that further complicate modeling (e.g., Gnat & Sternberg 2007). Generally, having access to more metal ion tracers means one is able to place more refined constraints on the models, whereas results from models with fewer ions are more model dependent.

Radiative transfer models like Cloudy (Ferland et al. 2013) are used to build PIE models (e.g., Bergeron & Stasińska 1986, Prochaska et al. 2004, Lehnert et al. 2013, Werk et al. 2014, Turner et al. 2015). These models are parameterized by density n_{H} , or equivalently the ionization parameter $\log U \equiv \Phi/n_{\text{H}}c$, the observed neutral gas column density N_{HI} , and a gas-phase metallicity, $\log [Z/\text{H}]$. Here, Φ is the number of photons at the Lyman edge (i.e., the number of ionizing photons), set by the assumed incident radiation field with a given flux of ionizing photons. Besides ionization and thermal equilibrium, another major underlying assumption of photoionization modeling is that the included metal ions arise from a single gas phase with the same origin (i.e., are cospatial). The single cloud, single density approximation for PIE modeling of low ions leads to uncertain cloud sizes, determined by $N_{\text{H}}/n_{\text{H}}$ ranging from 0.1 to 100 kpc (Stocke et al. 2013, Werk et al. 2014). In response, some models have begun to explore internal cloud density structure (Stern et al. 2016) or local sources of radiation (e.g., star formation in the galaxy, the hot ISM; Fox et al. 2005, Werk et al. 2016). PIE models generally fail for highly ionized metal species like OVI, sometimes CIV, and certainly for X-ray ions. For those we turn to CIE, where temperature controls the ionization fractions and a metallicity must be assumed or constrained to derive total hydrogen column N_{H} .

Beyond PIE and CIE, there are nonequilibrium ionization mechanisms that may reproduce the intermediate- and high-ion states that generally fail for PIE (e.g., CIV, NV, OVI). These models include: (a) radiative cooling flows that introduce gas dynamics and self-photoionization to CIE models (Edgar & Chevalier 1986, Benjamin 1994, Wakker et al. 2012); (b) turbulent mixing layers, in which cool clouds develop skins of warm gas in Kelvin–Helmholtz instabilities (Begelman & Fabian 1990, Slavin et al. 1993, Kwak & Shelton 2010); (c) conductive interfaces, in which cool clouds evaporate and hot gas condenses in the surface layer where electron collisions transport heat across the boundary (Gnat et al. 2010, Armillotta et al. 2016); and (d) ionized gas behind radiative shocks, perhaps produced by strong galactic winds (Dopita & Sutherland 1996, Heckman et al. 2002, Allen et al. 2008, Gnat & Sternberg 2009). These models all modify the column density ratios given by pure CIE but do not change the basic conclusion that gas bearing these ionic species must be highly ionized, i.e., with a neutral fraction $\ll 1\%$. These large and unavoidable ionization corrections, when applied to HI column densities of $\log N_{\text{HI}} \sim 15\text{--}18$, entail surface densities and total masses that are significant for the galactic budgets (Section 5). It is likely that combinations of PIE and CIE into these more complex models are more accurate descriptions of Nature than either basic process considered in isolation.

4.3. Line Profiles and Gas Kinematics

Line widths, given by the Doppler b parameter, illuminate the CGM temperature structure and gas dynamics. The gas temperature, T , and any internal nonthermal motions are captured in the following parameterization: $b^2 = (2kT/m_i) + b_{\text{nt}}^2$, for a species with atomic mass m_i . When the low and high ions are assessed via Voigt profile fitting, the low ions are usually consistent with gas temperatures $< 10^5$ K, with a contribution from nonthermal broadening (< 20 km s $^{-1}$; Tumlinson et al. 2013, Churchill et al. 2015, Werk et al. 2016). “Broad Lyman alpha” ($b \gtrsim 100$ km s $^{-1}$) and NeVIII systems have been detected in QSO spectra at high signal-to-noise ratios that directly probe gas at $\log T \sim 5.7$ (Narayanan et al. 2011, Savage et al. 2011, Tripp et al. 2011, Meiring et al.

Intermediate ions:

$40 \gtrsim \text{IP (eV)} \lesssim 100$,
 $T = 10^{4.5\text{--}5.5}$ K

PIE: photoionization equilibrium

CIE: collisional ionization equilibrium

High ions: $\text{IP} \gtrsim 100$ eV, $T > 10^{5.5}$ K

Low ions: $\text{IP} < 40$ eV,
 $T = 10^{4\text{--}4.5}$ K

2013). These UV absorption surveys indicate that the CGM contains a mixture of photoionized and/or collisionally ionized gas in a low-density medium at 10^4 – $10^{5.5}$ K (e.g., Adelberger et al. 2003, Richter et al. 2004, Fox et al. 2005, Narayanan et al. 2010, Matejek & Simcoe 2012, Stocke et al. 2013, Werk et al. 2013, Lehner et al. 2014, Savage et al. 2014, Turner et al. 2015).

The velocity dispersion and number of components reveal the kinematic substructure of the CGM. Most significantly, gas near low- z galaxies across the full range of $\log M_* = 8.5$ – 11.5 show projected line-of-sight velocity spreads that are less than the inferred halo escape velocity, even accounting for velocity projection. Thus most of the detected CGM absorption is consistent with being bound to the host galaxy, with implications for outflows and recycling (Section 7). This is true for all the observed species from HI (Tumlinson et al. 2013) to MgII (Bergeron & Boissé 1991, Johnson et al. 2015b, Nielsen et al. 2015) to OVI (Tumlinson et al. 2011, Mathes et al. 2014). The strongest absorption seen in HI and low ions is heavily concentrated within ± 100 km s $^{-1}$. For low ionization gas, internal turbulent/nonthermal motions are $b_{\text{nt}} \sim 20$ km s $^{-1}$, whereas for high ionization gas the nonthermal/turbulent contributions to the line widths are 50–75 km s $^{-1}$ (Werk et al. 2016, Faerman et al. 2017). Similar total line widths are seen in the $z > 2$ KODIAQ sample, possibly indicating similar physical origins at different epochs (Lehner et al. 2014).

Misalignments of the high and low ion absorption profiles in velocity space may indicate that the gas phases bearing high and low ions are not co-spatial and thus that the gas is multiphase (e.g., Fox et al. 2013). Some systems, however, show close alignment between low and high ionization gas (Tripp et al. 2011) in a fashion that suggests each detected cloud is itself multiphase, perhaps in a low ion cloud/high ion skin configuration. Heckman et al. (2002) and others (e.g., Grimes et al. 2009, Bordoloi et al. 2016) have argued that the relationship between OVI column density and absorption-line width for a wide range of physically diverse environments indicates a generic origin of OVI in collisionally ionized gas. However, the relationship exhibits considerable scatter, is impacted significantly by blending of multiple unresolved components [at least at the moderate $R \sim 20,000$ resolution of the Cosmic Origins Spectrograph (COS)], and may arise from other physical scenarios such as turbulent mixing (e.g., Tripp et al. 2008, Lehner et al. 2014). Generally, high ions like OVI in the CGM exhibit systematically broader line widths than low and intermediate ions (e.g., Werk et al. 2016). Though complex and varied, absorber kinematics may provide important observational constraints on both ionization and hydrodynamic modeling, but new methods of analysis and new statistical tools will be required to realize their full potential.

4.4. Challenges in Characterizing the Multiphase CGM

Ionization modeling is limited by what might be considered subgrid processes that investigators must cope with to get from line measurements to useful constraints on models. The most basic of these arise in the data themselves. CGM absorption observations are generally not photon-noise limited, but line saturation is a major issue particularly for the most commonly detected species. Only lower limits can be derived from the equivalent widths of saturated lines; line profile fitting helps where the saturation is not too severe. Reliable columns of the crucial HI ion are often challenging except where the Lyman limit is available. Moreover, the blending of narrow components with small velocity offsets in data with finite spatial resolution make all line measurements somewhat ambiguous. It is often necessary to model an entire line profile as a single nominal cloud, though sometimes the ionization state can be constrained on a component-by-component basis.

There is often ambiguity about whether to adopt PIE, CIE, or combination nonequilibrium models. These issues are compounded by uncertainties in the additional model inputs. These include the relative elemental abundances, which need not be solar but are usually assumed to

be. The extragalactic ultraviolet background (EUVB) is a particular problem as it may be uncertain especially at low redshift (Kollmeier et al. 2014), introducing up to an order of magnitude systematic error into some ionic abundances (Oppenheimer & Schaye 2013a).

EUVB: extragalactic
ultraviolet background

Though OVI is among the strongest and most frequently detected CGM metal absorption lines, it amply demonstrates the problems encountered in precisely constraining the exact physical origins of ionized gas. For example, absorption-line studies in high-resolution and high-S/N QSO spectra and complementary studies of HVCs around the MW show that the ionization mechanisms of OVI are both varied and complex over a wide range of environments (e.g., Sembach et al. 2004, Tripp et al. 2008, Savage et al. 2014). Ionic column density ratios and line profiles sometimes support a common photoionized origin for OVI, NV, and low-ion gas (e.g., Muzahid et al. 2015), whereas other systems require OVI to be collisionally ionized in a $\sim 10^{5.5}$ -K plasma (e.g., Tumlinson et al. 2005, Fox et al. 2009, Narayanan et al. 2011, Tripp et al. 2011, Wakker et al. 2012, Meiring et al. 2013, Turner et al. 2016). Often, the multiple components for a single absorber show both narrow and broad absorption lines consistent with both scenarios.

All these thorny issues with ionization modeling highlight the difficulty of getting at the detailed subgrid physics of a complex, dynamic, ionized medium. We should maintain a cautious posture toward conclusions that depend sensitively on exact ionization states. Much of the detailed physics is still at scales that we cannot yet resolve. Nevertheless, in Section 5, we see what we can learn by simplifying the situation to the most basic classes of models and proceeding from there.

4.5. Gastrophysical Models

The Galactic Corona began with Spitzer’s insight that cold clouds could be confined by a hot surrounding medium. This model has matured over the years into a strong line of theoretical research focused on the detailed physics of how the thermal, hydrodynamic, and ionization states of CGM gas evolve in dark matter halos. Placing multiphase gas into the context of the dark matter halo, Maller & Bullock (2004) suggested cold clouds cooled out of thermal instabilities in a hot medium while maintaining rough pressure equilibrium [though see Binney et al. (2009) for a counterpoint]. Accretion may also be seeded by gas ejected from the disk, as in the galactic fountain or precipitation model (e.g., Fraternali & Binney 2008, Voit et al. 2015b). These scenarios start with very simple assumptions—such as hydrostatic hot halos, diffuse clouds in PIE, or particular radial entropy profiles. These simplifying assumptions are necessary because we do not know the large-scale physical state of the CGM as a whole. Photoionization modeling of the low-ionization CGM using only the EUVB (Haardt & Madau 2001) strongly disfavors hydrostatic equilibrium with hot gas at T_{vir} (Werk et al. 2014); the cool and hot phases appear to have similar densities, rather than similar pressures. Furthermore, if OVI-traced gas follows a hydrostatic profile at the temperature where its ionization fraction peaks, $T \sim 10^{5.5}$ K, then its column density profile would be significantly steeper than observed (Tumlinson et al. 2011). There may be other means of supporting this gas, such as turbulence (Fielding et al. 2017), cosmic rays (Salem et al. 2016), or magnetic fields.

Adding to the uncertain physical conditions in the CGM is that OVI likely represents a massive reservoir of warm gas (Section 5.2.3). Such a massive reservoir is apparently at odds with the short cooling times for OVI given by typical CIE models; these timescales are often much shorter than the dynamical time, on the order of $\lesssim 10^8$ year. Yet, the short cooling times for OVI are in fact characteristic of many models for the multiphase CGM. In many formulations, the cooler low-ion-traced gas precipitates out of the warmer OVI-traced phase, owing to thermal instabilities (Shapiro & Field 1976, McCourt et al. 2012, Voit et al. 2015b, Thompson et al. 2016; see also Wang 1995), whereas the OVI-traced gas may be continually replenished by a hot galactic outflow.

In a similar vein, the OVI-traced warm gas could be cooling isochorically out of a hotter halo (e.g., Edgar & Chevalier 1986, Faerman et al. 2017) but overcome its short expected lifetime by extra energy injection from star formation or AGN.

Fully understanding the broader context and origin of the multiphase CGM requires more than what microphysical and phenomenological models alone can offer. Cosmological hydrodynamic simulations with self-consistent cosmic accretion and multiphase outflows are key to deciphering the panoply of observed absorption lines (Section 7). Moreover, much of the microphysics proposed as a natural source or maintainer of multiphase gas (e.g., thermal instabilities and turbulence) requires resolutions much higher than can be achieved by simulations that must simultaneously model the enormous dynamic range required for galactic assembly. Yet essentially all cosmological hydrodynamic simulations do produce a multiphase CGM (see, e.g., **Figure 3**). In general, the combination of the simulated density and temperature profiles of the CGM results in different ions preferentially residing at different galactocentric radii, with low ions preferring the denser, cooler inner CGM and higher ions filling the lower-density, hotter outer CGM (Hummels et al. 2013, Ford et al. 2014, Suresh et al. 2017; see also Stern et al. 2016). Yet inhomogeneous mixing of the different gas phases complicates predictions for gas cooling rates and the small-scale metal mixing that depend crucially on the unknown diffusion coefficient (Schaye et al. 2007).

Hydrodynamic simulations may be compared directly with observations via synthetic spectra, potentially helping to disentangle the degeneracy between physical space and observed velocity space. Constructing these synthetic spectra, however, faces many of the same challenges as modeling the ionization states of the observed gas: though the density, temperature, and metallicity of the simulated gas may be known, the EUVB and ionization mechanism must still be assumed in order to calculate ionization states (see, e.g., Hummels et al. 2016). Most simulations rely on the same radiative transfer codes (e.g., Cloudy; Ferland et al. 2013) that observational analyses do, though nonequilibrium chemistry and cooling are being included as computation power increases (Oppenheimer & Schaye 2013b, Silvia 2013). If these assumptions are incorrect, comparisons of derived results (such as masses) rather than observables (such as column densities) may lead to simulations getting the “right answer” for the wrong reasons.

5. THE BARYONIC MASS DISTRIBUTION OF THE CGM

5.1. The Missing Baryons Budget

Empirically constraining the total CGM mass as a function of stellar and/or halo mass is essential to quantifying models of galactic fueling and feedback. Under the condition $\Omega_b/\Omega_m = 0.16$ (Planck Collaboration et al. 2013), the total baryonic budget of sub- L^* to super- L^* galaxies spans two orders of magnitude, ranging from $10^{10.3}$ to $10^{12.3} M_\odot$. Although the stars and ISM for super- L^* galaxies are similar fractions of the total ($\sim 5\%$), the absolute amount of mass that must be found is around $100\times$ larger for sub- L^* galaxies and $10\times$ larger for L^* galaxies. How much of this 80–90% missing mass is in the CGM? We organize this subsection by temperature and review the observations, assumptions, and uncertainties in each calculation, using **Figures 7** and **8** to synthesize current results. We note that a recent review by Bland-Hawthorn & Gerhard (2016) performed a similar radially varying mass-budget compilation for the MW and its halo and incorporates some of these same results.

The baryon census as presented here relies on the assumption that galaxies fall along well-defined scaling relations of ISM and CGM gas mass as a function of stellar mass and that the scatter in these scaling relations is uncorrelated. We caution that there is tentative evidence that this is not necessarily the case: COS-GASS has shown galaxies with more cold gas in their ISMs




Figure 7

A synthesis of CGM mass density results for cold gas (*pink*, Zhu & Ménard 2013a), cool gas (*purple*, Werk et al. 2014), warm gas traced by OVI (*green*; Tumlinson et al. 2011, Peebles et al. 2014), X-ray-emitting gas (*yellow*, NGC1961; Anderson et al. 2016), and dust (*brown*, Ménard et al. 2010). An NFW profile scaled by Ω_b/Ω_m for $M_{\text{DM}} = 2 \times 10^{12} M_\odot$ is at the top in black. Abbreviations: CGM, circumgalactic medium; DM, dark matter; NFW, Navarro–Frenk–White.

have more cold gas in their CGMs (Borthakur et al. 2015). Although the correlation between CGM and ISM exhibits a high degree of scatter, likely from patchiness in the CGM, it exists at $>99.5\%$ confidence, and stacked Ly α profiles for low and high ISM masses clearly show the effect. The large-scale environment and gaseous interstellar content are difficult to explicitly account for in overall baryon budgets and may account for some of the scatter in the various estimates. For example, Burchett et al. (2016) find that the detection of CIV around galaxies with $M_\star > 10^{9.5} M_\odot$ drops significantly for galaxies in high-density regions (see also Johnson et al. 2015a). Future work should control for these properties.

5.2. CGM Masses by Phase

We now synthesize observations from across the electromagnetic spectrum to estimate CGM masses and mass profiles for cold, cool, warm, and hot gas. These results are plotted together in **Figures 7 and 8**.

5.2.1. Cold gas, $T < 10^4$ K. Cold gas tracers consist of neutral and low ions like HI, NaI, CaII, and dust. This is material that may have cooled from hotter phases that experienced thermal

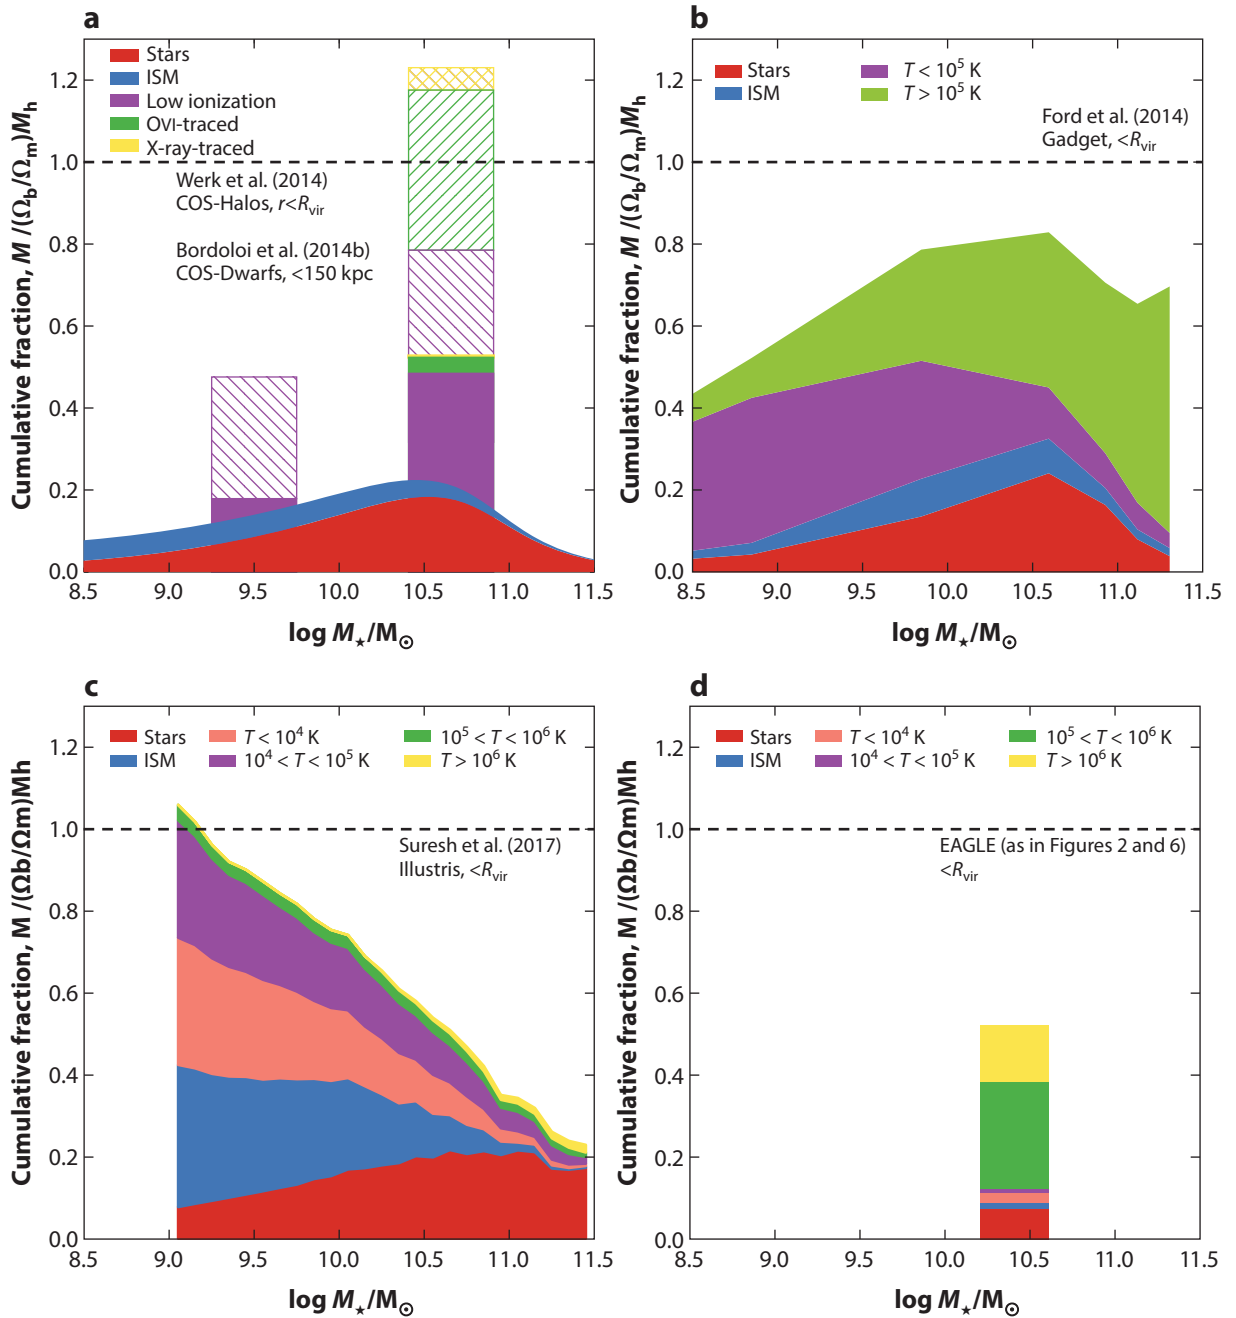


Figure 8

(a) An accounting of CGM baryon budgets for all physical phases. The solid bars show the minimum values, whereas the hatched regions show the maximal values. The other three panels show simulated baryon budgets from (b) Ford et al. (2014), (c) Illustris (Suresh et al. 2017), and (d) the EAGLE halo shown in **Figures 2 and 6** (Schaye et al. 2015, Oppenheimer et al. 2016). Abbreviations: CGM, circumgalactic medium; ISM, interstellar medium.

instability or may arise in clouds entrained in multiphase outflows. Putman et al. (2012) estimated the total cold gas mass traced by HVCs in the MW halo to be $M = 2.6 \times 10^7 M_\odot$ (including only HVCs detected via 21-cm emission, and excluding the Magellanic Stream system). The Magellanic Stream provides an additional contribution of $\sim 3 \times 10^8 M_\odot$, but it cannot be assumed to be a generic feature of galaxies. Thus, the total contribution from cold gas is $M \lesssim 10^9 M_\odot$ even if the interstellar gas of the Magellanic Clouds is included, making up less than 1% of the missing baryons for an MW-like halo. We further note that though dust masses have been estimated from stacks of reddened background QSOs (Ménard et al. 2010) and galaxies as standard crayons (Peek et al. 2015), indicating values comparable with the dust in the ISM of these galaxies (see Section 6.3), both ISM and CGM dust are at most only $\sim 1\%$ of the missing halo baryons. Finally, using stacked optical spectra from SDSS, Zhu et al. (2013) derived a column density profile for gas bearing CaII H and K around $\sim L^*$ galaxies. For the purposes of **Figure 5**, we have converted this to a mass density profile, conservatively assuming that the calcium is entirely in CaII and $Z = Z_\odot$. The total mass for CaII itself is $5,000 M_\odot$, and when we scale to $Z = Z_\odot$, we derive $M = 2 \times 10^8 M_\odot$ for the cold component, again $\sim 1\%$ or less of the baryon budgets.

5.2.2. UV Absorption lines and the cool 10^{4-5} -K CGM. The mass of the cool CGM ($\sim 10^{4-5}$ K) is perhaps the best constrained of all the phases at low redshift, owing to the rich set of UV lines in this temperature range. Prior to COS, estimates for this phase were based on single ions with very simple ionization and metallicity corrections to arrive at rough estimates. Prochaska et al. (2011a) estimated $M_{\text{cool}} \approx 3 \times 10^{10} M_\odot$ for all galaxies from $0.01 L^*$ to L^* , assuming a constant $N_{\text{H}} = 10^{19} \text{ cm}^{-2}$ out to 300 kpc. Using a blind sample of MgII absorbers, Chen et al. (2010) estimated $M_{\text{cool}} \approx 6 \times 10^9 M_\odot$ for the MgII-bearing clouds alone. The former estimate simply took a characteristic ionization correction, whereas the latter counted velocity components as clouds and converted from a metal column density to N_{H} using a metallicity, because neither study had the multiphase diagnostic line sets that could be used to self-consistently constrain gas density and metallicity. Both L^* and super- L^* galaxies have provided the most reliable constraints, mainly because of their relative ease of detection in photometric and spectroscopic surveys at $z < 0.5$ (Chen & Mulchaey 2009, Prochaska et al. 2011a, Werk et al. 2012, Stocke et al. 2013).

With COS, it became practical to build statistically significant samples of absorbers that cover a broader range of ions. These estimates still rely on photoionization modeling, carried out under the standard assumption that the low ions and H I trace cool ($T < 10^5$ K) gas, and the primary source of ionizing radiation is the extragalactic UV background. Using the COS-Halos survey, Werk et al. (2014) addressed the mass density profile and total mass for $L \approx L^*$ galaxies with PIE models that derive self-consistent n_{H} and Z using a range of adjacent ionization states of low-ion absorption lines (primarily CII, CIII, SiII, SiIII, NII, and NIII). The resulting surface density profile appears in **Figure 7** and yields $M_{\text{cool}} = 6.5 \times 10^{10} M_\odot$ for L^* galaxies out to R_{vir} . Using the same COS-Halos sample with new COS spectra covering the Lyman limit, and taking a nonparametric approach with a robust treatment of uncertainties, Prochaska et al. (2017) recently refined the cool CGM mass estimate to be $9.2 \pm 4.3 \times 10^{10} M_\odot$ out to 160 kpc. Stocke et al. (2013) used a complementary approach to estimate individual cloud sizes and masses, along with their average volume filling factor, for galaxies in three luminosity bins ($< 0.1 L^*$, $0.1-1 L^*$, and $L > L^*$). They find volume filling factors that range from 3% to 5% for their modeled clouds, with length scales ($N_{\text{H}}/n_{\text{H}}$) in the range of 0.1–30 kpc, totaling $\log M_{\text{cool}} = 7.8-8.3, 9.5-9.9$, and $10-10.4$, respectively. Finally, Stern et al. (2016) determine the total mass in the cool (and possibly warm CGM) of $1.3 \pm 0.4 \times 10^{10} M_\odot$ for L^* galaxies given their universal cloud density profile. In this phenomenological model each ion occupies a shell of a given n and T such that the fraction of gas

in that particular ionization state is maximized. Thus, this calculation represents a conservative minimum of baryons that must be present. These ranges are shown in **Figure 8**.

For super- L^* galaxies, Zhu et al. (2014) use stacking techniques to estimate the correlation function between luminous red galaxies with a mean stellar mass of $10^{11.5} M_\odot$ and cool gas traced by MgII absorption in SDSS data for $\sim 850,000$ galaxies with $0.4 < z < 0.75$. The cool CGM around massive galaxies calculated in this way appears to completely close the CGM baryon budget for super- L^* galaxies, at 17% of the total halo mass. The assumptions for metallicity and ionization corrections, however, make it uncertain.

5.2.3. UV Absorption lines and the warm 10^{5-6} -K CGM. In **Figure 6**, ions like CIV, NV, OVI, and NeVII trace the warm CGM at $T \approx 10^{5-6}$ K. However, this temperature range in particular is burdened by significant uncertainty in the precise ionization mechanism responsible for its purported ionic tracers (see Section 4.4). If high ions are partially photoionized, OVI, for example, may trace a non-negligible fraction of $T < 10^5$ K gas that has already been counted toward the total baryon census in the previous section. For gas traced by OVI, Werk et al. (2016) point out that typical photoionization models like those used for the low ions have difficulty accounting for the total column of OVI and column density ratios of NV/OVI without the need for path lengths in excess of 100 kpc. However, significant additional ionizing radiation at ~ 100 eV may reduce this requirement.

In general, CIE models require a very narrow range of temperature to reproduce the OVI observations, $T = 10^{5.3-5.6}$ K (Tumlinson et al. 2011, Werk et al. 2016). Furthermore, the kinematics of OVI relative to the low ions, in particular large b values, seem to naturally support the idea that the OVI is in a hotter phase (Tripp et al. 2011, Muzahid et al. 2012; see also Tripp et al. 2001, Stern et al. 2016). Tumlinson et al. (2011) found that OVI traces a warm CGM component that contributes $> 2 \times 10^9 M_\odot$ of gas to the L^* baryon budget. This mass estimate is strictly a lower limit due to the conservative assumptions adopted: (a) solar metallicity; (b) the maximum fraction of oxygen in OVI allowed by CIE models, 0.2; and (c) the CGM sharply ends at 150 kpc. We adopt $\log M_{\text{warm}} = 10.0$ in **Figure 8** for the COS-Halos galaxies (see also Faerman et al. 2017).

For sub- L^* galaxies, Bordoloi et al. (2014b) estimate M_{warm} using CIV. These galaxies are at $z < 0.1$, and the COS spectra do not cover the full range of Lyman series lines and ions available at $z > 0.1$, hindering detailed ionization modeling. COS only covers OVI at $z > 0.2$, where it is difficult to assemble statistically significant samples of confirmed sub- L^* galaxies, so an OVI-based mass estimate for low-mass galaxies is not currently possible. With these caveats in mind, assuming a limiting ionization fraction for CIV, Bordoloi et al. derive $\log M_{\text{warm}} = 9.5$ if the gas typically has solar metallicity. For gas with lower metallicity, e.g., 0.1 solar, the value is 10 times higher and rather closer to baryonic closure for sub- L^* galaxies (**Figure 8**). We caution that for CIV, detailed photoionization often places CIV with low-ionization state gas rather than with high-ionization state gas (e.g., Narayanan et al. 2011). Thus, the CIV-derived mass for sub- L^* galaxies is highly uncertain without detections of additional ionization states.

One of the most surprising results to emerge from Tumlinson et al. (2011) is that OVI appears to be absent around the non-star-forming, more massive galaxies in the COS-Halos sample. Thus, there is tentative evidence that $\sim 10^{5.5}$ K gas is not a major component of the CGM of super- L^* galaxies, which may be a result of massive galaxies having generally hotter halos or nonequilibrium cooling (Oppenheimer et al. 2016). Thus, we do not have a good observational constraint for the warm CGM baryonic content for super- L^* galaxies. The EUV ion NeVIII redshifts into the COS band at $z > 0.5$, where a few detections (Tripp et al. 2011, Meiring et al. 2013) hint that it may be present in halos out to 100–200 kpc. However, the number of absorbers associated with particular galaxies is not yet sufficient to include it in mass estimates for the warm phase.

5.2.4. The hot $T > 10^6$ -K phase. Hot gas at the virial temperature ($T_{\text{vir}} = GM_{\text{halo}}m_p/kR_{\text{vir}}$) is a long-standing prediction. For $M_{\text{halo}} \gtrsim 10^{12} M_{\odot}$, the temperature should be $T \gtrsim 10^6$ K, and observable at X-ray wavelengths, although there are EUV tracers such as MgX and SiXII that have yet to yield positive detections (**Figure 4**). Only a few very luminous spirals and ellipticals have had their halos detected (Anderson & Bregman 2011, Dai et al. 2012, Bogdán et al. 2013, Walker et al. 2015, Anderson et al. 2016), and independent constraints on the temperature, density, and metallicity profiles from soft X-ray spectroscopy are rarer still. Thus the fraction of baryons residing in the hot phase, and its dependence on stellar and or halo mass, are not yet determined.

Three sets of constraints are relevant: the MW, individual external galaxies, and stacked samples of external galaxies. Anderson & Bregman (2010) addressed directly the problem of whether hot gas could close the baryon budget for the MW. From indirect constraints such as pulsar dispersion measures toward the Large Magellanic Cloud, cold gas cloud morphology, and the diffuse X-ray background, they limited the hot gas mass to $M \lesssim 0.5\text{--}1.5 \times 10^{10} M_{\odot}$ or only 2–5% of the missing mass. The choice of an NFW profile for the hot gas is a key assumption: If the density profile is assumed to be flatter ($\beta \sim 0.5$), the mass can be 3–5 times higher, but still only 6–13% of the missing baryons. The claims by Gupta et al. (2012) that the baryon budget is closed for the MW, based on the assumption of an isothermal, uniform density medium, have been questioned by evidence that the gas is neither isothermal nor of uniform density (Wang & Yao 2012).

The well-studied case of NGC 1961 (Anderson et al. 2016) constrains the hot gas surface density out to $R \simeq 40$ kpc, inside which $M_{\text{hot}} = 7 \times 10^9 M_{\odot}$. This is only 2% of the stellar mass ($3 \times 10^{11} M_{\odot}$) and so leaves the galaxy far from baryonic closure. Extrapolating to 400 kpc yields $M_{\text{hot}} = 4 \times 10^{11} M_{\odot}$, but given the declining temperature profile it is likely that the gas well beyond the X-ray detections is at lower temperatures, $T \lesssim 10^6$ K, where EUV and FUV indicators provide the best diagnostics. Stacked emission maps of nearby galaxies provide the strongest evidence for extended hot halos. In a stack of 2,165 isolated, K-selected galaxies from ROSAT (*Röntgensatellit*), Anderson et al. (2013) found strong evidence for X-ray emission around early-type galaxies and extremely luminous galaxies of both early and late type. The X-ray luminosity depends more on galaxy luminosity than on morphological type. Luminous galaxies show $M = 4 \times 10^9 M_{\odot}$ within 50 kpc, and $M = 1.5\text{--}3.3 \times 10^{10} M_{\odot}$ if extrapolated out to 200 kpc, which is comparable with the stellar masses. Yet high amounts of hot gas this far out would appear to be excluded, as found by Yao et al. (2010), who stacked *Chandra* spectra at the redshifts of foreground galaxies and placed strict ($\lesssim 1$ mÅ) limits on OVII and OVIII. The limits are also consistent with the limits on nearby galaxy emissivity, derived earlier by Anderson & Bregman (2010). The key uncertainty is how far out the hot gas extends with the flat, $\beta \sim 0.5$ density profile seen at $R \lesssim 50$ kpc, but the Yao et al. (2010) limits imply that hot gas halos around nearby galaxies host at most $\simeq 10^{10} M_{\odot}$. In their summary of the X-ray results, Werk et al. (2014) adopted $M_{\text{hot}} = 1\text{--}14 \times 10^9 M_{\odot}$ from Anderson et al. (2013).

The thermal Sunyaev–Zel’dovich (SZ) effect—scattering of CMB photons by free electrons in a plasma—may constrain the hot gas content of galaxy clusters and halos down to the galactic scale. The Planck Collaboration et al. (2013) and Greco et al. (2015) claim detections down to $M_{\star} = 2 \times 10^{11} M_{\odot}$ and a possible signal down to $M_{\star} = 6 \times 10^{10} M_{\odot}$. These results create tension with the X-ray measurements, because the SZ detections imply a self-similar relation between M_{halo} and M_{hot} down from the cluster scale ($M_{\text{halo}} \sim 10^{14} M_{\odot}$), where we know hot baryons close the budgets, into the galactic range where this is much less clear. It may be that the hot gas extends well beyond the X-ray surface brightness limits at 50 kpc up to the megaparsec scales where the SZ effect is measured. However, if every $\geq L^*$ halo was filled with T_{vir} gas, it would violate constraints from the soft X-ray background (Wu et al. 2001). If halos with $M_{\text{halo}} \lesssim 10^{11} M_{\odot}$ depart from self-similarity, it could be the cooling and feedback that prevent halos from reaching

CMB:

cosmic microwave background

SZ:

Sunyaev–Zel’dovich

their cosmic share of baryons. The kinematic SZ effect—in which photons receive a Doppler shift when scattering from a plasma with bulk motion—may be valid at even lower masses for halo gas measurements (Hill et al. 2016). This work is in its early stages, and we look forward to more progress that complements studies in the UV and X-ray bands.

5.2.5. Theoretical considerations. From the discussion above and the synthesis in **Figure 8a**, we see that CGM measurements have added significantly to the baryon budgets for galaxies and may complete those budgets under some assumptions. There has been theoretical progress as well: Hydrodynamical simulations generally agree that the CGM contains a budget of baryons at the same order of magnitude as the stellar masses. In **Figure 8b–d**, we show there is less quantitative agreement for the temperature partitioning of the CGM as a function of stellar mass, despite these models having approximately the same predictions for the baryonic content of galaxies.

A promising aspect of this quantitative disagreement is that different physical treatments of energetic and/or kinetic feedback do indeed lead to different total baryon fractions, and in particular they lead to different trends in the fraction by phase. Thus, observations of how CGM gas masses are distributed by phase can favor or disfavor particular physical prescriptions and thus already offer phenomenological tests of models. However, these comparisons additionally show how challenging it will be to perform stringent tests. Even where simulations with radically different physical prescriptions yield opposite trends, at any particular mass they are only different by factors of $\lesssim 2$ in the fraction of any phase. At present, this range is comparable with the systematic errors remaining in the observational characterization of the phases. Thus any claims that the data favors or disfavors any particular model should be made and interpreted carefully. As discussed in Section 4.5, comparing the models with observations—by using synthetic data and directly comparing observables such as column densities and line kinematics—has the benefit of shifting the myriad assumptions discussed in Section 4.2 onto the simulations.

6. METALS: NATURE’S TRACER PARTICLES

6.1. The Metals Census

Total mass budgets by themselves do not fully reveal the flows that govern galaxy evolution. However, there is a ready means of distinguishing inflows from outflows: Stars produce heavy elements sending passively advecting tracer particles out into the ISM, CGM, and IGM from stellar winds and supernovae. The metal content of galactic flows can help identify their origins and determine their fate, and it can break degeneracies between models matched to the four galaxy problems. The galactic metals census (Section 2) requires that we compute the total budget of available metals produced by the galaxy by $z = 0$. This census was performed by Peebles et al. (2014) by compiling measurements on stars, ISM and CGM gas, and dust. As shown in **Figure 9**, the contributions bound in stars, interstellar gas, and interstellar dust—the metals inside galaxies—add up to only consistently 20–30% over a factor of $\sim 1,000$ in stellar mass. (Though the overall level of the fraction of metals retained in galaxies is uncertain, primarily owing to uncertainties in nucleosynthetic yields, the flatness of this relation is fairly robust; see Peebles et al. (2014) for a thorough discussion of the uncertainties in this calculation.) Ideally, this census would be done for each element individually, with the CGM divided into each ionization state of that element, e.g., oxygen (Oppenheimer et al. 2016), but as that is observationally not yet generally feasible, the ionization corrections discussed in earlier sections must instead be done to account for unobservable ionization states. Qualitatively similar results are seen in simulations that have addressed this problem in particular (Muratov et al. 2017). This striking invariance must offer

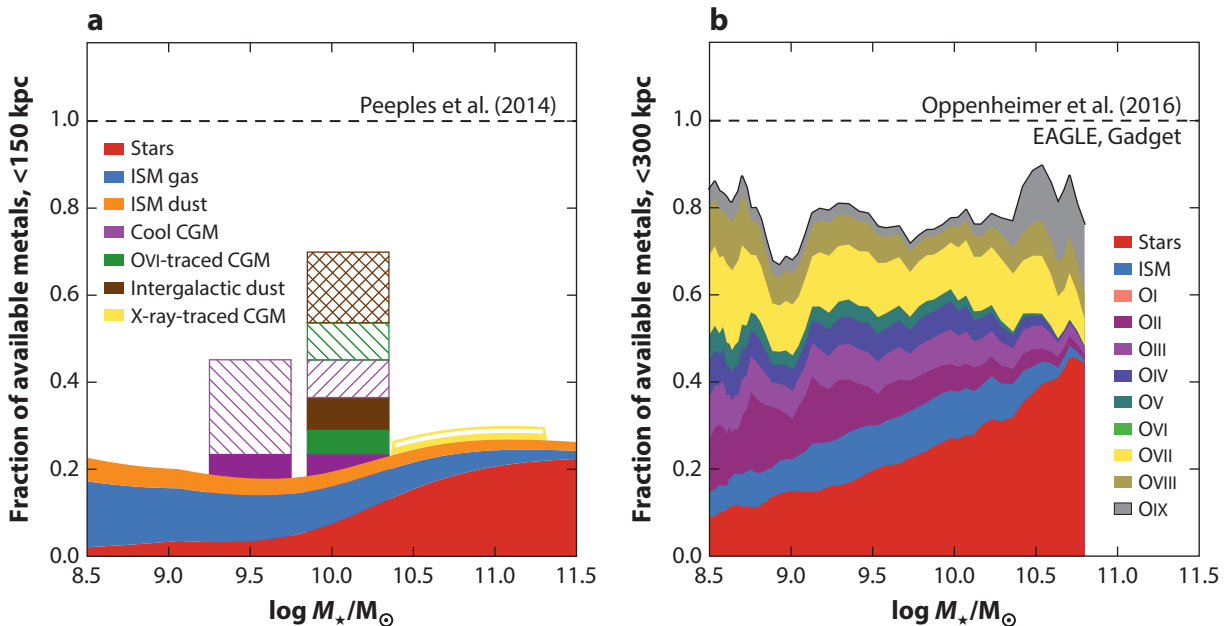


Figure 9

(a) A metals census of the CGM around star-forming $z \sim 0$ galaxies following Peebles et al. (2014), including a sub- L^* budget from Bordoloi et al. (2014b). As in Figure 7, stars are red, ISM gas is blue, ISM dust is orange, the cool CGM is purple, the OVI-traced CGM is green, the X-ray-traced CGM is yellow, and intergalactic dust is brown. (b) A simulated budget from 55 relatively isolated $\log M_* \geq 8.5$ star-forming EAGLE halos, with a moving average smoothing (Oppenheimer et al. 2016). In both panels, the denominator is the total mass of metals ever produced by the central galaxy; the CGM may have contributions from, e.g., satellites. Abbreviations: CGM, circumgalactic medium; ISM, interstellar medium.

some important clues to the operation of galactic outflows and inflows, with potentially large implications for the processes of galaxy fueling, feedback, and recycling.

6.2. Metals Observed as Gas

Even Lyman Spitzer might have recognized that the heavy elements observed in the CGM are in some sense the cause of, and solution to, all our problems. Apart from the (problematic) series of Lyman lines in the rest-frame FUV, virtually all our knowledge of the physical state, mass, kinematics, and evolution of the CGM gas come from lines of C, N, O, Si, Fe, Mg, Ca, and so on, whether they appear in the UV or X-ray bands. Yet, as described in Section 4, these critical diagnostics also present many problems of analysis and interpretation. To work through this, it helps to distinguish between measurements of metal content or metal mass, on one hand, and metallicity on the other. This distinction hinges on whether or not the hydrogen content can be measured, which is notoriously difficult. Measurements of hydrogen suffer severe H I saturation effects, and juggling both metals and hydrogen compounds the difficulties of ionization corrections. When considering metal mass, we can often tolerate simpler ionization corrections or even direct sums of metal ion surface densities, sidestepping the large ionization corrections for H I (Section 4.4).

The COS-Halos survey (Tumlinson et al. 2011) used the OVI line observed with COS in a way that typifies measurements of metal content rather than metallicity. Their basic empirical finding is that OVI appears at column densities of $\log N_{\text{OVI}} \simeq 14\text{--}14.5$ out to the 150-kpc limits

of the survey. Because OVI does not reach more than 20% of the total oxygen in most ionization conditions, they were able to place a robust lower limit of $>10^7 M_{\odot}$ of total oxygen for star-forming galaxies. As it comes from direct integration of surface densities for a heavy element, does not refer to H, and uses a limiting ionization correction, this estimate avoids some of the trickiest aspects of metallicity measurements and yet has significant implications for the budgets of galactic metals (Peeples et al. 2014). The OVI traces a high ionization component of the CGM gas; adding lower ionization gas to the budget requires the more complex ionization corrections and assumed relative abundances of oxygen and, e.g., Mg and Si, though it does not require the HI-dependent metallicity corrections that plague the baryon census. Altogether, 20–30% of available metals have been located in the $R < 150$ kpc CGM around $\sim L^*$ galaxies.

By contrast with the measurements of total metal mass, bona fide metallicities require robust measurements of the hydrogen surface density, which entails accurate measurements of N_{HI} and reliable ionization corrections. For most strong CGM absorbers at $z \lesssim 0.2$, the Lyman series lines are saturated and do not yield reliable HI column densities. However, beyond this redshift, and at $\log N_{\text{HI}} > 16.2$, LLSs enable adequately precise (± 0.2 – 0.3) measurements of N_{HI} , and the ionization corrections are manageable.

By building a sample of LLSs from high-quality COS sightlines, Lehner et al. (2013) and Wotta et al. (2016) found that the distribution of metals in LLSs clearly exhibits two peaks near 4% solar and 50% solar metallicity (**Figure 10a**). The metallicities are constrained by detections of low to intermediate ions such as CII–CIV, SiII–SiIV, OII–OIV, and MgII. This bimodal distribution qualitatively matches with expectations that accretion from the IGM into halos will have low metallicity, whereas accretion of gas previously ejected will have higher metallicity. The relative absence of intermediate values challenges our intuition that gas should naturally mix over time into a continuous distribution and has posed a challenge to simulations (Hafen et al. 2016, etc.). But most of these systems have not yet been identified with galaxies. In contrast to the Lehner bimodality, Prochaska et al. (2017) find a unimodal distribution of metallicities within 160 kpc

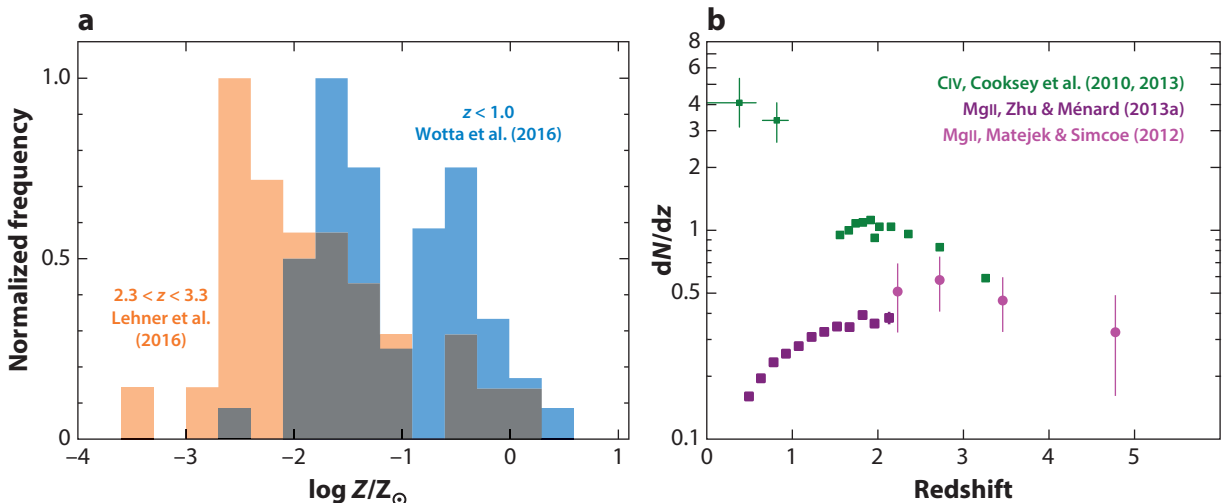


Figure 10

Two views of CGM metallicity. (a) Two LLS distributions from Lehner et al. (2013) and Wotta et al. (2016). This comparison clearly shows evolution in the LLS metallicities over time. (b) Trends in MgII and CIV line densities per unit redshift: the low-ion MgII traces the cosmic star-formation history, whereas CIV continually becomes more abundant. Abbreviation: LLS, Lyman limit system.

of L^* galaxies with a median of $\sim 30\%$ solar metallicity. These metallicities derive from tight constraints on N_{HI} around L^* COS-Halos galaxies with well-defined masses and distance to the absorber. The contrast between the absorber-selected Lehner et al. sample and the galaxy-selected COS-Halos sample may indicate that they arise in other selection effects, but it may also indicate variation in CGM metallicity in different subsets of the galaxy population.

By mining the Keck database of high- z QSO absorbers, the KODIAQ (Keck Observatory Database of Ionized Absorption toward Quasars) survey studied a sample of LLSs at $z > 2$ (Lehner et al. 2016). This sample is shown in **Figure 10a** compared with the expanded low- z sample of Wotta et al. (2016). The $z > 2$ distribution is unimodal and centered at $[X/H] \sim -2$. A similar result was obtained for two samples of LLSs at still higher redshift, $z = 3.5\text{--}4$, with unimodal distributions centered at $[X/H] \sim -2.5$ (Cooper et al. 2015, Glidden et al. 2016). This is near the bottom edge of the low-metallicity peak at $z < 1$, indicating evolution in the average metallicity of high-column CGM over the few gigayear interval. Somehow, the bimodality emerges long after the initial buildup of metals and is noticeable only in the $z < 1$ sample. Note that neither of these samples has specific galaxies attached—both are selected based on HI alone, and the galaxies will have to be identified later. It is also possible that the column density range used for selection traces different galaxy masses, radii, and total column densities at the different redshifts, and so the apparent evolution does not occur in the same type of physical system (owing to a higher mean cosmic density). Nevertheless, it is now possible to compare the distribution of CGM metallicities over $\sim 6\text{--}10$ Gyr of cosmic time.

In particular, there are ever-increasing samples of $z > 2$ absorbers that do have associated galaxy information, allowing for a more direct comparison with the low- z COS studies (**Figure 4**). The Keck Baryonic Structure Survey (Rudie et al. 2012) has engaged in a long campaign to characterize the CGM of star-forming galaxies at $z \sim 2.2$, going back to pioneering studies of absorption associated with Lyman-break galaxies (Adelberger et al. 2003). These data show ion sets that overlap strongly with the low- z studies. Both HI and metals (OIV, NV, CIII, CIV, and OVI) show strong statistical correlations with galaxies out to 100–300 kpc. Using stacking, Steidel et al. (2010) and Turner et al. (2015) examined the relative kinematics of metals and galaxies, finding essentially only outflow kinematics and little sign of inflow; there must be gas flowing in to maintain the observed SFRs, but it may be occurring in thin filaments with low covering fraction.

These results across redshift can be viewed a different way, by examining the redshift evolution of strong lines that are likely to trace CGM gas. **Figure 10b** shows the comoving sightline densities of MgII ($W_{\text{rest}} \geq 1 \text{ \AA}$) and CIV ($W_{\text{rest}} \geq 100 \text{ m\AA}$), which follow different trends at $z < 2$. The number density of strong MgII absorbers rises and then declines again toward $z = 0$. Absorbers above this limit occur within ~ 100 kpc of galaxies (see **Figure 4**), so the resemblance of this curve to the cosmic SFR density (Hopkins et al. 2006) suggests that the strong MgII absorbers are linked to the fueling or feedback of star formation. Indeed, other evidence suggests that we are seeing the rise and decline of galactic superwinds (see Section 7.3). In contrast to the MgII, strong CIV absorbers continue their march upward at low redshift. This trend in moderate-to-high ionization gas may indicate that ionized gas in occupying the bulk of the CGM volume becomes more common even as strong winds creating MgII absorbers decline with the cosmic SFR density.

6.3. Metals Observed as Dust

The ISM is a mixture of gas and dust; this is no less true of the CGM. In a pioneering study, York et al. (2006) stacked a sample of 800 strong MgII absorbers to find evidence of Small Magellanic Cloud-like dust reddening. Ménard et al. (2010) added the SDSS photometric galaxy catalogs to this style of analysis and found that the reddening extends over angular scales consistent with

distances hundreds of kiloparsecs away from the luminous galaxies (**Figure 7**). To tie dust to specific galaxies and precise physical scales, Peek et al. (2015) used passively evolving galaxies from SDSS as standard crayons to examine the reddening imposed by foreground SDSS spectroscopic galaxies. They found a strong reddening effect out to 150 kpc in the bluest bands and a steeper drop past that radius than that in the angular correlations of Ménard et al. (2010). The correlations with physical radius allow Peek et al. (2015) to further estimate the typical total mass of dust for galaxies between 0.1 and $1L^*$ of $M_{\text{dust}} \simeq 6 \pm 2 \times 10^7 M_{\odot}$. They found only a weak trend with stellar mass, $M_{\text{dust}} \propto M_{\star}^{0.2}$, and no discernible trend with the galaxies’ specific SFRs. Thus the presence of dust in the CGM out to 100-kpc scales provides unambiguous evidence that the CGM is fed by galactic outflows, accounting for approximately 10% of the metals budget near L^* (**Figure 9**). This degree of reddening can be explained by outflows from normal star-forming galaxies in simulations, provided the dust-to-gas ratio is similar to the Galactic value and the dust survives the trip (Zu et al. 2011). It is not yet clear why the dust properties show so little dependence on galaxy stellar mass, resembling the CGM H I and low ions more than the CGM high-ionization gas. It might be that the increasing reddening at low redshift indicates a steady buildup of metals in the CGM and a relative lack of recycling into future star formation. Dust observations could also be used to test the physical models of galactic outflows that employ radiation pressure on dust to drag gas out of galaxies (Murray et al. 2005, 2011). Further explorations of CGM dust promise to constrain galactic outflows and recycling in ways that complement studies of gas.

7. INFLOWS, OUTFLOWS, AND RECYCLING

7.1. The Problems: Galaxy Fueling and “Missing” Metals

Recent findings show that the CGM possesses a significant budget of baryons, but how are they feeding galaxies across the spectrum of galaxy masses (**Figure 2**)? Accreting gas passes through the CGM on its journey from the IGM to galaxies, where it presumably leaves some observable signatures that we can use to characterize the inflows. The rates of accretion onto galaxies and of outflow out of galaxies are crucial parameters in most models of galaxy evolution (Tinsley 1980). However, there is not agreement about where and how a galaxy’s fuel source is regulated. It is often assumed gas inflow from the IGM is balanced by the sum of star formation, gas ejection as outflows, and any net buildup of gas in the ISM (Lilly et al. 2013, Dekel & Mandelker 2014, Somerville & Davé 2015). This formulation completely omits the role of the CGM, even at the phenomenological level, but this “bathtub” model appears to nonetheless describe the many broad trends in galaxy scaling relations with redshift (Dekel & Mandelker 2014). These models, though they do not explicitly address the CGM’s composition or physical state, nonetheless have specific implications for its content and evolution (e.g., Shattow et al. 2015). Conversely, models that use physical principles to describe the regulation of flows between the CGM and ISM (Voit et al. 2015b) can reproduce the same phenomenological galaxy scaling relations without detailed treatments of star formation inside galaxies. With observations of the CGM and its dynamics, we can potentially assess whether its role in regulating star formation is trivial, as the former models assume, or essential, as the latter models assume. Ideally, CGM observations would not only answer this question but also reveal how it fuels star formation and manages outflows as a function of galaxy mass.

The observations we have discussed up to this point reveal the CGM (at low z) as a massive gaseous medium with a rich internal kinematic structure that is, in bulk, consistent with being bound to the host galaxies. Yet the degeneracy between kinematics and the physical location of absorbing gas can easily get lost in transverse sightline observations.

In simulations the CGM can appear to have obvious and well-ordered large-scale structure, with accreting and outflowing gas occupying physically distinct regions such as filaments and biconical outflows (Shen et al. 2012, Corlies & Schiminovich 2016; see also **Figure 3**), but at low redshift, circumgalactic gas tends to be more well mixed, with instantaneous velocities having little bearing on the origin or fate of a particular pocket of gas (Ford et al. 2014, Muratov et al. 2015, Christensen et al. 2016), though this is also seen at $z = 3$ (van de Voort et al. 2012). In light of the observational projection effects, and theoretical cautions, we now consider what can be learned from observing inflow and outflow directly in down-the-barrel observations, in which we interpret gas blueshifted relative to the galaxy as outflowing and redshifted gas as inflowing. These observations are better at probing gas in or near the disk-halo interface rather than the proper CGM out in the halo. Considering them in conjunction with CGM findings from transverse sightlines promises insights into the dynamics of the CGM that are not otherwise available.

7.2. Empirical Signs of Fueling and Inflows

Gas accretion is perhaps the most fundamental process in galaxy formation (Fox & Davé 2017), as galaxies must acquire gas while feedback is optional. In the prevailing theoretical paradigm, gas flowing into galaxies at $\lesssim 10^{12} M_{\odot}$ should be dynamically and thermally cold, whereas more massive halos receive most of their baryons as hotter ($T > 10^5$) gas (Dekel & Woo 2003, Kereš et al. 2005, Dekel & Birnboim 2006, Kereš & Hernquist 2009, Stewart et al. 2011, though see Nelson et al. 2013). Thus cold, dense, metal-poor CGM gas is often interpreted as direct evidence of accretion. First, cool, dense CGM gas is abundant in the form of LLSSs. A large fraction of these are metal poor at all redshifts (Lehner et al. 2013, Glidden et al. 2016, Cooper et al. 2015). Metal-poor LLSSs are evident as tracers of accretion in high-resolution simulations (Fumagalli et al. 2011, Hafen et al. 2016). The cool, bound HI seen in the CGM of $z \sim 0.25$ galaxies (Tumlinson et al. 2013) should have a short cooling time. Finally, the finding from COS-GASS that there is a correlation between interstellar and circumgalactic HI (Borthakur et al. 2015) implies a connection between circumgalactic fuel and star-forming fuel. Though sub- L^* and dwarf galaxies have not yet had their cool CGM masses measured directly, the widespread presence of Ly α at similar strength suggests they too possess significant budgets of cold halo gas.

All this evidence taken together strongly indicates that galaxies possess large reservoirs of CGM gas eligible for accretion. Yet evidence for fuel does not automatically constitute evidence for fueling: bound, cold gas has turned up in halos in which its presence is surprising, such as the CGM of passive galaxies (Thom et al. 2012). The actual fate of this material is unclear: How can we claim the bound cold gas is fueling star-forming galaxies but not the passive galaxies? We therefore seek direct signatures of gas accretion onto galaxies. Yet these signatures are notoriously difficult to observe because incoming material may be metal poor, ionized, and obscured by outflowing material. Once gas is near the disk, proving empirically that it is accreting can be extremely difficult when it is seen in projection and its kinematics are easily confused with disk material.

The MW itself provides direct and unambiguous evidence for inflow in the form of its blueshifted HVCs and the striking Magellanic Stream. The HVCs arise in many complexes of clouds lying within ~ 10 kpc of the disk and have $100\text{--}300 \text{ km s}^{-1}$ blueshifted radial velocities that indicate they will reach the disk within $10^{7\text{--}8}$ years. Their mass inflow rate falls in the range of $0.1\text{--}0.5 M_{\odot} \text{ year}^{-1}$, compared with the $1\text{--}2 M_{\odot} \text{ year}^{-1}$ of star formation (Putman et al. 2012). These clouds are all detectable in 21 cm emission, meaning that they occupy the tip of the column density distribution of CGM gas seen around other galaxies. The inflow rate inferred for ionized gas is much larger than that for the classical HVCs, $\dot{M} \simeq 0.8\text{--}1.4 M_{\odot} \text{ year}^{-1}$ (Lehner & Howk 2011), which is more comparable with the MW's SFR. The Magellanic Stream is estimated to

contain around $2 \times 10^9 M_\odot$ of gas in neutral and ionized form (Bland-Hawthorn et al. 2007, Tepper-García et al. 2015), and could provide $\sim 5 M_\odot \text{ year}^{-1}$ of gas to the MW disk as it accretes (Fox et al. 2014). Unfortunately, HVCs both above and below the radio-detection threshold are difficult to detect in external galaxies, despite intensive searches (Putman et al. 2012), and satellites like the Magellanic Clouds and their Stream are not very common in L^* galaxies. So, we cannot generalize this result to mainstream galaxy populations.

Down-the-barrel spectroscopy provides complementary information on inflows. Using this technique on $z \sim 0.5$ galaxies with Keck spectroscopy and *Hubble* imaging, Rubin et al. (2012) detected clear signs of inflow at $80\text{--}200 \text{ km s}^{-1}$ in star-forming galaxies of $\log M_*/M_\odot = 9.5\text{--}10.5$, inferring mass inflow rates of $\dot{M} \gtrsim 0.2\text{--}3 M_\odot \text{ year}^{-1}$. It seems likely that these estimates significantly undercount inflow, because inflowing (redshifted) gas is often obscured by outflows (blueshifted) or by emission from the galaxy’s ISM (this problem is especially noticeable at higher redshift; Steidel et al. 2010). Even if outflow is not present, the profiles are not sensitive to accretion from the lower half of the bimodal LLS metallicity distribution (Lehner et al. 2013), which could make up a large fraction of the available cold CGM gas. Recently, Zheng et al. (2017) reported the detection of enriched, accreting gas at the disk-halo interface of M33 via COS observations of SiIV absorption along several sightlines to bright O stars in the disk. Their kinematic modeling of the observed absorption features implies an accretion rate of $2.9 M_\odot \text{ year}^{-1}$. Although these results provide evidence for accretion of cold, metal-enriched gas directly into galaxy disks, evidence for more metal-poor cold-mode accretion, and for gas entering farther out in the disk (on-ramp; **Figure 1**), is still lacking (though see Bouché et al. 2013), as is empirical characterization of how accretion rates vary with galaxy mass.

7.3. The Preeminence of Outflows

By consensus, outflows are an accomplice if not the perpetrator in each of the problems outlined in Section 2. The existence of outflows is not in question: The large share of metals outside galaxies provides incontrovertible evidence for them (Section 6). COS-Halos found widespread OVI around star-forming galaxies—extended to $\sim 300 \text{ kpc}$ by Johnson et al. (2015b)—but could not show that this ion becomes more prevalent with SFR. Even so, simulations found that robust outflows were necessary to produce the observed reservoir of metals (e.g., Ford et al. 2013, Hummels et al. 2013, Suresh et al. 2017), such that the high-metal ions provide a significant constraint on the time-integrated effects of outflows even if it does not show the effects of recent or ongoing outflows directly. After that, the important questions concern how they transport baryons, metals, momentum, energy, and angular momentum. There is empirical evidence and strong theoretical suggestions that the physical drivers and properties of galaxy winds—their velocity, mass loading, metal content, and likelihood of escape—depend on galaxy mass, circular velocity (v_{circ}), SFR, and metallicity. Many investigators pursue CGM observations in the hope that they can help to constrain these outflows and how they scale with galaxy properties.

Direct observational evidence for outflows is readily available at all redshifts [see Veilleux et al. (2005) for a review]. In the nearby Universe, large-scale complex multiphase outflows are seen in starbursts (e.g., M82) and from the MW’s central regions (Fox et al. 2015). Down-the-barrel spectroscopy of the NaI D in local starbursts (Martin 2005) found that outflow velocities depend linearly on v_{circ} . Rubin et al. (2012) and Bordoloi et al. (2014a) characterized similar flows using MgII at $z \sim 1$. At $z > 2$, where the FUV-band ions used at $z \sim 0$ appear at visible wavelengths, Steidel et al. (2010) used down-the-barrel spectroscopy to detect nearly ubiquitous outflows in rapidly star-forming Lyman break galaxies, with no clear indications for redshifted inflow. Although these results help constrain the mass loading and covering fraction of outflows,

they do not show how far these winds propagate into the CGM. It may be that the bulk of the energy is transported out in the hot gas while the bulk of the mass leaves in the cold phase, but this is still an open question (Strickland & Heckman 2009).

Absorbers on transverse sightlines can directly constrain the impact of winds on the CGM. Cross-correlations of MgII absorbers with the orientation of galaxies on the sky at $z \lesssim 1$, from both samples of individual galaxies (Kacprzak et al. 2012, Mathes et al. 2014) and stacked spectroscopy (Bordoloi et al. 2011, Zhu & Ménard 2013a), find that the strongest absorbers prefer the semiminor axis of disk galaxies, as expected for biconical outflows emerging from the disk. The preference for the semiminor axis disappears by $\sim 60\text{--}80$ kpc, indicating that winds propagate at least that far or merge into the general medium near that radius (e.g., the $z = 2$ example in **Figure 1**). Studies of outflow covering fractions at $z \sim 1$ reinforce a picture of outflows being roughly biconical, with little surface area ($\sim 5\%$) solely dedicated to inflow (Martin et al. 2012, Rubin et al. 2014). Another strong clue about outflows comes from examining the CGM of starburst and post-starburst galaxies. Using an SDSS-selected sample, Heckman & Borthakur (2016) found unusually strong HI and multiphase ions at 100–200 kpc compared with the COS-Halos and COS-GASS samples of galaxies at lower SFRs. These studies collectively show that SFR is a factor in determining the content of the CGM, perhaps as far out as R_{vir} .

Down-the-barrel measurements tell us that outflows are ubiquitous, and sightline measurements tell us that they reach 100 kpc scales. Together these findings suggest that a large part of the CGM is made of outflows, and to examine one is to illuminate the other. The open questions concern not only the basic scaling of velocity and mass loading with galaxy v_{circ} —which has received much attention—but just as importantly the distribution of outflow temperatures, metallicities, and fate. These cannot (yet) be simulated from first principles but can be constrained by the combination of CGM and down-the-barrel observations. The former constrain the radial extent and the velocity fields of multiphase gas far from the disk, whereas the latter constrain the initial velocities, mass loading, and (possibly) metallicities.

A recent goal of models and simulations has been to discriminate between winds that are momentum-driven (Murray et al. 2005), which appear to improve the match of simulations to the galaxy mass–metallicity relation (Finlator & Davé 2008) and the metal content of the IGM (Oppenheimer & Davé 2006, 2008), and those that are energy-driven (Murray et al. 2011), which appear to better match the galaxy stellar mass function (Davé et al. 2012) and new COS data (Ford et al. 2016). A momentum-driven outflow has a velocity $v_w \propto v_{\text{circ}}^{-1}$, whereas an energy-driven flow has much faster outflows for low-mass galaxies with $v_w \propto v_{\text{circ}}^{-2}$; with a fiducial wind speed of $\sim 100 \text{ km s}^{-1}$, an unimpeded flow reaches 100 kpc in only 1 Gyr, i.e., the scales on which metals are seen in the CGM (Section 6). Thus understanding the history of CGM metals and the velocities and mass flow rates of galactic flows go hand in hand. Real winds may depend less on the local potential well and more on the local SFR surface density (Kornei et al. 2012, Heckman et al. 2015). New hydrodynamic simulations of galaxies that resolve the multiphase ISM and explicitly include radiation pressure and thermal pressure (Hopkins et al. 2012) support this picture. Like essentially every other simulation suite on the market, however, models with this feedback scheme have too little OVI in the CGM while retaining too many metals in stars (Muratov et al. 2015).

7.4. Following the Metals: The Role of Recycling

Inflow and outflow are necessary processes in galaxy and CGM evolution; can one become the other by the recycling of outflows into fresh accretion of ejected gas? We have already established that, at least at low redshift, galaxies require a long-term source of fuel and that their CGM gas and metals are massive and bound. Recycling is a natural consequence; this gas should reaccrete

onto the galaxy if the cooling time is short. Indeed, the predominance of metal-enriched accretion is supported by essentially all cosmological simulations where the origins of gas joining the ISM has been tracked: Significant fractions of gas accreting onto galaxies has previously been ISM gas—and often through multiple cycles (Ford et al. 2014, Christensen et al. 2016, Muratov et al. 2017), with the majority of star formation at late times being fueled by recycled gas (Oppenheimer et al. 2010). Ford et al. (2014) found 60% of all star formation at $z = 0$ is powered by gas that was in the CGM a billion years before. This idea has the intriguing implication that a substantial fraction of all heavy elements on Earth once cycled through the MW’s halo at 100-kpc scales. The timescales are unclear: Christensen et al. (2016) find that half of outflow mass is recycled on a timescale of 1 Gyr with a logarithmic tail, independent of halo mass, whereas Oppenheimer & Davé (2008) find that $t_{\text{rec}} \propto M_{\text{halo}}^{-1/2} \sim 10^{9 \pm 0.5}$ years, which is a timescale so short for massive galaxies that it is like not having an outflow at all and so long for dwarfs that it essentially escapes forever.

Thus the idea of recycling is well motivated, but the details are still murky. Is it a simple process in which gas launched at $v < v_{\text{esc}}$ encounters hydrodynamic resistance and eventually succumbs to gravity to fall back into the galaxy as part of a large-scale halo fountain? Or is the CGM well mixed but multiphase, with metal-rich gas precipitating out of the hot halo and raining onto the galaxy (Fraternali et al. 2015, Voit et al. 2015a, Thompson et al. 2016)? Here, too, metals can help disentangle the ins and outs. Intriguingly, dense CGM gas (Lehner et al. 2013, Wotta et al. 2016; Section 6.2) is roughly equally divided between gas at a few percent solar (metal-poor IGM accretion) and 40% solar (recycling ejecta?).

Although gas accreting from the IGM generally has (or is assumed to have) very low metallicity (Lehnert et al. 2013, Cooper et al. 2015, Glidden et al. 2016), cases with metallicity well below the IGM ($\text{Ly}\alpha$ forest) at the same redshift are rare (Fumagalli et al. 2011, Crighton et al. 2017). That is, pristine cosmic accretion either entrains metal-enriched circumgalactic gas on its way into the galaxy (e.g., Fraternali et al. 2015) or, even at the highest redshifts where accretion is potentially observable, is at least partially comprised by material that has previously been in the ISM; i.e., recycled mode accretion is critical to galaxy evolution even at early cosmic times. Yet most formulations of the bathtub model assume that the accreting gas is pristine [e.g., Lu et al. (2015), though see Davé et al. (2012)]. Entrainment is a commonly invoked phenomenon for galaxy outflows; it refers to the wind fluid sweeping up ambient ISM and mixing it with the fresh supernova ejecta powering the outflow. [It is important to note that the metallicity of the outflowing material is necessarily higher than that of the ambient ISM, contrary to what is assumed in some popular simulation recipes (e.g., Vogelsberger et al. 2014).] Does recycled accretion behave in a similar way but in the opposite direction, with pristine inflows sweeping up metal-polluted CGM material on its way from the IGM to the ISM? Or do galaxy winds preferentially reaccrete, sweeping up more pristine cosmic accretion?

Taking all this evidence into account, we can see the outlines of an emerging picture of galaxy inflows, at least at low redshift. They arise in the massive reservoir of cold, metal-enriched gas bound to a galaxy’s potential well, and enter the disk in HVC-like clouds but also in smooth flows of ionized gas. There may be a metal-poor component that comes more directly from the IGM without spending much time in the CGM or otherwise acquiring metals. All these aspects of the CGM—cold, bound, metal enriched, and accreting—align better with the phenomenon of recycled accretion than the bimodal hot/cold accretion. Recycled accretion arises from the ejection of metal-enriched galactic winds that lack the energy to escape the halo entirely or which encounter the CGM itself and lose energy to radiation from shocks and then eventually cool and re-enter the galaxy. It may be that recycling, rather than accretion and feedback, is the more accurate way of viewing how galaxies acquire their gas.

8. THE PARADOX OF QUENCHING

Passive and/or quenched galaxies possess little if any cold gas in their ISM, and blaming the CGM merely relocates the problem: How and why do these massive galaxies that once possessed a cold ISM lose and not regain it? Presumably their dark matter halos continue to add mass, but the accompanying gas does not enter the ISM and form stars like it once did. How galaxies achieve this transition is a deep and abiding problem in astrophysics, and the array of possible mechanisms for consuming, removing, and/or heating cold gas are beyond the scope of our review. We address the phenomenon of quenching by considering the CGM as a factor in, and indicator of, the quenching process.

LRG: luminous red galaxies

8.1. The Fate of Cold Accretion and the Problem with Recycling

The accretion of gas into halos, its heating to around T_{vir} , and its eventual cooling and entry to the ISM was long the prevailing picture of galaxy fueling. In an important twist on this basic picture, Kereš et al. (2005) argued that star-forming galaxies are fed by “cold accretion”: the gas never reaches T_{vir} but enters a galaxy’s disk via streams while remaining below $T \sim 10^5$ K. Above $\log M_*/M_\odot \sim 10.3\text{--}10.5$ (or $M_{\text{halo}} \sim 10^{12} M_\odot$), the dark matter halo has sufficient mass, and the CGM enough pressure, to support a virial shock and suppress the cold mode. The coincidence of this mass with the stellar mass that divides star-forming from passive galaxies (**Figure 2**) drew great attention to this scenario (e.g., Dekel & Birnboim 2006), leading to predictions that the halos of passive galaxies should possess little cold gas (Stewart et al. 2011).

The observational picture belies the clean transition seen in simulations and the stark division of observed star-formation rates. Although COS-Halos did find a dramatic difference in highly ionized OVI around star-forming and passive galaxies, the latter do not show as strong a deficit of CGM HI. As shown in **Figure 11**, the equivalent widths and covering fractions of HI do not drop as stellar mass increases across the range $\log M_* \simeq 10\text{--}11$ (Thom et al. 2012). This is directly contrary to the expectation from, e.g., Stewart et al. (2011) that the covering fraction of strong HI should drop to nearly zero as galaxies transition to the hot mode of accretion. The inner CGM (<50 kpc), however, is not well covered by these observations (**Figure 4**); it is possible that high-pressure hot gas close to the galaxy prevents this cold material from accreting, as some models predict (Schawinski et al. 2014).

The presence of cool gas in the halos of massive red galaxies is now well established by MgII studies. Gauthier et al. (2010) and Bowen & Chelouche (2011) found covering fraction of $f_c = 10\text{--}20\%$ out to 100–200 kpc for >1 Å absorbers around luminous red galaxies (LRGs). Using a sample of $\sim 4,000$ foreground galaxies at $z = 0.5\text{--}0.9$ from the zCOSMOS survey, Bordoloi et al. (2011) found that the MgII equivalent width for blue galaxies is 8–10 times stronger at inner radii (<50 kpc) than for red galaxies, but even red galaxies possess evidence for cold gas. Using a new SDSS-based catalog of MgII QSO absorbers and LRGs, Zhu et al. (2014) mapped the mean profile out to $\gg 1$ -Mpc scales and argue that the mean profile at this mass scale is even stronger than that found by Bordoloi et al., extending at a detectable level out to 1 Mpc for LRGs. Johnson et al. (2015b) have pointed out that strong MgII absorbers are usually consistent with being bound to their host halos, meaning that the cold gas is contained with the dynamical influence of the galaxy.

From a theoretical perspective, the quenching of galaxies is still a significant unsolved problem. Star formation must be curtailed, and later accretion and cooling of gas must be suppressed indefinitely to explain how galaxies remain passive for >6 Gyr (Gallazzi et al. 2008). Theories vary in how they accomplish this: Some models artificially truncate star formation based on halo mass (Somerville & Davé 2015), whereas others suppress the star-forming fuel by heating the CGM itself (e.g., Gabor et al. 2010, Gabor & Davé 2012). Thus the CGM itself can be the proximate

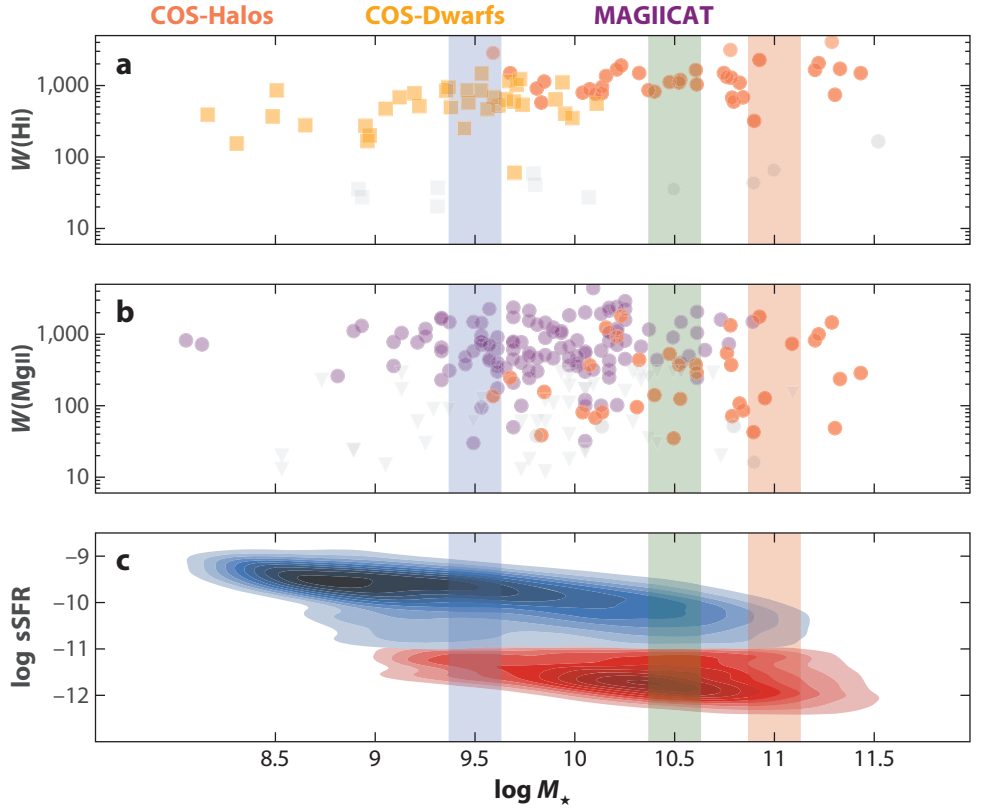


Figure 11

Three views of the CGM and quenching. (a) A trend in Ly α equivalent width over three decades in stellar mass from COS-Halos (Tumlinson et al. 2013) and COS-Dwarfs (Bordoloi et al. 2014a). As shown by Thom et al. (2012), the presence of HI around red, passive galaxies indicates that their halos are not devoid of cold gas. (b) MgII from COS-Halos and MAGIICAT (Nielsen et al. 2016). Gray points indicate upper limits. (c) The galaxy SFR bimodality from **Figure 2**. Abbreviations: CGM, circumgalactic medium; SFR, star-formation rate.

cause of quenching, even if the source of CGM heating is not yet identified. Unfortunately models that manipulate the CGM directly cannot be tested against CGM observations, or at least, they must be modified somehow to recover the cold gas seen in passive galaxy halos.

By contrast, models that include self-consistent subgrid treatments of feedback, whether thermal (Schaye et al. 2015), mechanical (Choi et al. 2015), or a combination of thermal, mechanical, and radiative (Vogelsberger et al. 2014), can be compared with CGM observations as tests of their success. As an example, the mechanical feedback model implemented by Choi et al. (2015) performed better than the standard (Springel et al. 2005, Di Matteo et al. 2008) thermal feedback model in both suppressing galaxy formation and reducing the surface density of gas in the CGM by factors of 3–10 at 10–100 kpc.

Suresh et al. (2017) addressed quenching using the Illustris simulations, which are tuned to the observed M_*/M_{halo} and galaxy metallicities but not the CGM. In Illustris, thermal AGN feedback is deposited locally, inside the galaxy, when the supermassive black hole is in its energetic

quasar mode. But in the $\sim 90\%$ of the time when the supermassive black hole is accreting quiescently, its radio mode feedback is deposited nonlocally as thermal energy over 100-kpc scales. This amounts to direct heating of the CGM, shifting cold gas to intermediate temperatures showing more OVI and otherwise warm gas to high temperatures showing OVII and OVIII. The net effect is that the Tumlinson et al. (2011) trend of strong OVI around star-forming galaxies and weak OVI around passive galaxies is recovered. The cold CGM is reduced but not completely destroyed. To be consistent, any visible effects of feedback would need to persist even when the AGN is not active, as the COS-Halos galaxies in question are not AGN at the time we observe them. The EAGLE simulations presented by Oppenheimer et al. (2016) show a similar conclusion with models of thermal feedback and nonequilibrium cooling: At higher mass, with more feedback, OVI is suppressed and the cold gas is depleted but not completely destroyed. These feedback effects force behaviors that generally resemble the data: They suppress star formation to create a red sequence, they force net gas loss from the inner CGM by heating gas that then buoyantly rises, and they shift the balance of gas ionization toward higher temperatures and higher ions.

Despite these advances, the basic paradox of quenching remains: What happens to the halos of passive galaxies to quench their star formation, keep it quenched, and yet leave cold gas present in their halos? If passive galaxies possess cold gas and are not using it, can we be sure of the (naïvely obvious) conclusion that star-forming galaxies are using the diffuse gas they possess? Moreover, if the bulk of star formation at low z comes from recycled accretion, then to understand both how galaxies get their gas and how galaxies quench, we must understand how both the internal and external fuel supplies are shut off.

8.2. The CGM of AGNs and Quasars

If feedback from AGNs is effective at quenching their star formation and their cold CGM in simulations, it naturally suggests that this effect will be visible in the gaseous halos of galaxies with ongoing AGN activity. Although hard radiation fields of AGNs may leave distinctive ionization signatures in halo gas even long after the AGN fades (Keel et al. 2012, Oppenheimer & Schaye 2013b), studies like COS-Halos with subsamples of passive galaxies have excluded active AGNs for the most part and even so have not seen any apparent signs of AGN effects on the CGM. No published study has systematically examined background QSO/foreground AGN pairs, though there is one such study underway with *Hubble*/COS (<http://www.stsci.edu/cgi-bin/get-proposal-info?id=13774>).

At $z > 2$, the Quasars Probing Quasars program has seen clear evidence that galaxies hosting bright quasars show greatly enhanced gas budgets in HI and low ions (Prochaska et al. 2014) though less excess in the high ions. This enhancement of neutral and low-ionization gas hints at a larger accretion rate for these robustly star-forming galaxies. AGNs may even yield a net gain of cold gas in the CGM (Faucher-Giguere et al. 2016). The Ly α blobs observed at $z > 2$ may be gas accreting on to galaxies, with radiation powered by gravitational infall (Goerdt et al. 2010), though these data may be more consistent with illumination from buried AGNs (Prescott et al. 2015). The higher gas masses only exacerbate the problem of feedback and quenching—there is more gas to be removed, and it is still not clear how that gas is removed or heated and accretion suppressed thereafter. Future work should focus on following such galaxies down through cosmic time as their QSOs fade, star formation is quenched, and the galaxies later evolve passively. Post-AGN and post-starburst galaxies should be examined for CGM gas as much as is practical. Understanding this process is critical to properly understanding the role of the CGM in creating or reflecting the birth of the red sequence. Refer to the sidebar titled Some Questions and Directions for Future Work for a summary of what we need to know about the red sequence.

SOME QUESTIONS AND DIRECTIONS FOR FUTURE WORK

Data in Need of More Theory

1. Are there any clean observational tests or theoretical discriminants between the various heuristic models of feedback?
2. Are there self-consistent models of quenching that produce a red sequence of galaxies and yet leave a significant mass of cold CGM? How is the remaining cold gas kept from accreting?
3. What do the detailed kinematic profiles of the multiphase suite of absorbing ions tell us about the physical and dynamic structures of the CGM?

Theory in Need of More Data

1. What is the mass and composition of the CGM at high redshift and in low- z $M_* < 10^{10} M_\odot$ galaxies, and how do these constrain galaxy evolution models?
2. What is the small-scale density and kinematic structure of the CGM, and what does it tell us about the physics?
3. What does the CGM do as galaxies quench? Does cool, neutral gas extend into the inner CGM of passive galaxies?
4. Where are the metals that are still missing from the census? What are the elemental abundance ratios in CGM gas, and how do they depend on the galaxy's mass and star-formation history?

9. OPEN PROBLEMS, FUTURE PROSPECTS, AND FINAL THOUGHTS

9.1. Progress and Problems

New instruments and new thinking reveal the CGM as a complex, dynamic gaseous environment that may close galactic baryon budgets and regulate gas accretion, star formation, and chemical enrichment. The observational studies that underlie the mass density profiles in **Figure 7** and mass budgets in Sections 4 and 5 have all been obtained since 2010. For years, questions about how and when gaseous halos influenced galaxy evolution consistently struggled with what was there. The bulk contents of the CGM are now better characterized than ever before. There remain missing pieces—the baryon and metals budget well below L^* must still be determined (**Figure 8**), and many of the metals are missing—but we can already see signs that the most urgent questions motivating new studies take what and where as known and go on to ask how and when. These sorts of questions strike more directly at physics than at phenomenology.

9.1.1. The scale problem. How a gaseous halo evolves is determined at any instant primarily by its density, temperature, metallicity, and radiation fields. But for an actual CGM (such as the one simulated in **Figure 6**), these physical quantities vary and evolve on many relevant scales, ranging from the subparsec sizes for single cold clouds to the more than 100-kpc size of the whole CGM and even more than megaparsec scales in the IGM. If we are to answer the hows of accretion, feedback, recycling, and quenching, we must achieve a better understanding of the basic physical fields at higher spatial and kinematic resolution. This means finding ways to capture subparsec boundary layers and instabilities while also maintaining the much-greater-than-kiloparsec context. Yet this 5–6 order-of-magnitude range still cannot be captured simultaneously in numerical simulations. One approach would be to continue the development of physically rigorous analytic models (e.g., Voit et al. 2015b, Fielding et al. 2017, Thompson et al. 2016, Faerman et al. 2017) that can isolate the key physical effects and then to incorporate these lessons into simulations at the subgrid level while their resolution improves with computing power. For instance, it might be possible to

include subgrid models that account for unresolved interfaces between hot and cold gas or to extract subgrid models for cosmological boxes from extremely high-resolution idealized cloud simulations with carefully controlled physics. To complete the leap between phenomenology and physics, these intrinsically subgrid processes must come under control while the proper cosmological and galactic context is maintained.

The transport of metals and the information they provide would also benefit from addressing the scale problem. Metals trace feedback and drive cooling, so how they are distributed through CGM gas at small scales is a critical factor in a proper physical understanding of accretion and feedback. Dense CGM gas appears to be bimodal in metallicity, congruent with the ideas of pristine accretion and recycled winds. What does this tell us about the small-scale structure of the CGM, the relationship between accretion and feedback, and the mixing of diffuse gas? These are among the thorniest of open questions, because of the huge dynamic range in metallicity that must be captured. This problem will be addressed by larger absorber and galaxy surveys but perhaps poses its stiffest challenges to numerical simulations, because many of the relevant physical mechanisms for mixing gas at boundaries and interfaces are still well below the subgrid level of simulations. This is another case in which coupling small-scale simulations of clouds to cosmological boxes could pay dividends.

The scale problem exists also for data but might be better labeled a problem of resolution and confusion. In data, the rich multiphase and multiscale structures of CGM gas are seen through a complex rendering in absorption or emission lines from diagnostic ions. The line profiles of absorbers likely contain more information than we are currently able to extract and interpret. Systematic effects from line saturation, uncertain ionization and radiation fields, relative abundances, limited signal-to-noise ratios, and finite spectral resolution all complicate the derivation of the true CGM density field, which in turn enters into mass estimates, energy balance, and timescales for the gas flows of interest. Although we are learning to model and simulate the CGM at higher resolution with better physics, we should also aim to extract and use the full information available in the rich kinematic profiles of multiphase absorbers, which will likely require new analytic and statistical techniques. The importance and complexity of the CGM make it imperative to examine all of the information that Nature provides.

9.1.2. Mass flows and the fate problem. The CGM matters to galaxies as long as it provides them with fuel and recycles their feedback. Ultimately, this is what we care about most—how does the CGM influence galaxy evolution? The most fundamental questions with which we began are still not completely answered: How does cold gas accrete and form stars over billions of years, and why does this cycle stop in massive galaxies? Does the CGM empty out or get consumed when galaxies quench? How much star formation is fueled by recycling and how much by new accretion? Can we ever hope to identify particular absorbers as accretion, feedback, or recycling, or are we destined never to separate them? These questions will drive the field as it advances from phenomenology toward more sophisticated physical understanding. Properly explaining these phenomena in terms of the hows of accretion, feedback, recycling, and quenching requires that we follow mass flows, not merely mass budgets.

Now that we have a grip on the bulk contents of the CGM, it is time to develop and deploy the tools to probe these questions of how the gas flows operate. To follow flows, we must make at least three key advances. First, the mass budgets should be characterized more fully in all phases at stages of galaxy evolution, including those that are relatively short lived such as mergers and AGNs. These analyses would additionally benefit from analyzing how outflows and inflows seen in down-the-barrel measurements relate to the kinematics viewed on transverse sightlines, an overdue synthesis deserving attention from both observations and theory. Second, we must

attempt to directly constrain the timescales of CGM evolution using data alone—how do mass budgets and kinematics jointly constrain timescales? Third, we must look at simulations in a new way that focuses on the origins and evolution of the physical phases and how these appear in the data. A large measure of simulation work addressed to the CGM has focused on using column densities and kinematics to constrain uncertain mechanisms of feedback by matching real data to mocks from simulations. While these issues are being resolved, it is also valuable to look at simulations from a different phenomenological point of view. The study from Ford et al. (2014) provides an example; their paper identified particles as “pristine accretion,” “recycled accretion,” “young outflows,” and “ancient outflows” and followed their evolution over time. These insightful categories turn out to be correlated with observable signatures. We believe there is great potential in viewing models and data from this angle, trying to identify the more distinctive or even unique manifestations of key physical processes defined by their fate rather than their instantaneous properties or appearance.

9.2. Future Prospects for Data

The next decade should bring a wide array of new instruments and numerical capabilities that will address these unsolved problems. While *Hubble* lasts (through the mid-2020s), UV absorber samples will grow, particularly those that focus on the $z > 0.5$ regime where a broader set of EUV ionization diagnostics is available (such as NeVIII). This increase in coverage will in turn allow more careful treatments of ionization diagnostics component by component, hopefully with a better understanding of how CGM gas is spread across physical phases and across galaxy mass. COS remains the ideal instrument for this problem, and big advances are still possible in the metals budget, ionization and kinematic relationships of multiphase gas, and the relationships between CGM gas and special types of galaxies. Starting in 2018, the *James Webb Space Telescope* will enable much deeper searches for faint galaxies near QSO sightlines, likely associating galaxies with samples of $z > 4$ absorbers that are already known (Matejek & Simcoe 2012, Becker et al. 2015). Detections of HI emission (e.g., Cantalupo et al. 2014, Arrigoni Battaia et al. 2015, Martin et al. 2015) will provide useful tests of models for CGM mass and structure, but the problems of gas ionization state and metal transport will require much more challenging maps of emission from oxygen and carbon ions [see Hayes et al. (2016) for a pioneering effort]. Such maps might emerge from integral field unit spectrographs such as MUSE (Multi Unit Spectroscopic Explorer) and KCWI (Keck Cosmic Web Imager), and their successors on 30-m-class telescopes; limits can be further improved by stacking of multiple galaxies. The optimal galaxies would be those where absorption-line probes are also available, so that emission-line and pencil-beam measurements can be compared. Emission maps of metal-bearing CGM gas (e.g., Bertone et al. 2010, Corlies & Schiminovich 2016) are a key goal of the Large Ultraviolet/Optical/Near Infrared Surveyor (LUVOIR; <http://asd.gsfc.nasa.gov/luvoir/>), which will push to $50\times$ the UV point source sensitivity of *Hubble*/COS and to 100-fold multiplexing in UV spectroscopy. Planned for launch in the 2030s, LUVOIR would be able to directly image the CGM in metal-line emission, map the most diffuse gas with weak absorbers, and resolve the multiphase kinematics of CGM gas with $R > 50,000$ UV spectroscopy (Dalcanton et al. 2015). The hot gas phase would be addressed by the ESA-planned X-ray flagship known as the *Advanced Telescope for High ENergy Astrophysics* (ATHENA; <http://sci.esa.int/cosmic-vision/54517-athena/>) in 2028, with a significant focus on understanding the cosmic evolution of hot gas in the IGM and CGM.

The size of our samples provides statistical power over the key galaxy variables: mass, redshift, shape, evolutionary state, and orientation to the sightline. Here, future UV absorber samples must be supplemented by optical absorber samples at $z \sim 3$ and by deeper galaxy surveys at all redshifts.

This is a problem for the next generation of giant ground-based telescopes, which will advance high- z CGM studies in rest-UV lines and support low- z studies by obtaining redshifts of sub- L^* galaxies near QSO sightlines at surveys at $z < 1$ to fill in the low-mass baryon and metals census, which is still a major missing piece.

Massive fiber-based surveys have proven effective at characterizing CGM gas and its flows with both intervening and down-the-barrel measurements. This technique should only accelerate in the future, pushing to fainter sources, higher redshifts, and rarer foreground galaxies with future massively multiplexed spectrographs (e.g., eBOSS, PFS) on large telescopes. This technique excels at detecting weak signals in the CGM and at examining more and more foreground galaxy properties with good statistics. With larger, deeper samples, we can look forward to addressing questions about the behavior of the cold/dense CGM in rarer galaxy types, such as quasars and AGNs, mergers, and groups.

9.3. Final Thoughts

Galaxies were understood as island universes long before astronomers discovered the interstellar gas that forms their stars. The IGM was added to the big picture with the discovery of QSO absorption lines and the development of the dark-matter cosmology. Because it is much fainter than stars, and much smaller than the IGM, the CGM is arguably the last major component of galaxies to be added, but it has nevertheless become a vital frontier. As to why, it is clear that much has been learned by viewing galaxy evolution from the perspective of the CGM. The CGM can even provoke fascination: Might the heavy elements on Earth be cycled back and forth through the MW's CGM multiple times before the formation of the Solar System? It appears that the solution to major problems in galaxy formation that are still unsolved will run through this elusive region of the cosmos.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

M.S.P. and J.T. acknowledge support from NSF grant AST-1517908. We are grateful to Ann Feild of Space Telescope Science Institute for her expert artistic contributions; Joop Schaye and Ben Oppenheimer for use of the EAGLE simulation shown in **Figures 2, 6, and 8**; Josh Suresh for data from the Illustris simulation shown in **Figure 8**; Sasha Muratov for data from the FIRE simulation (**Figure 8**); and Ben Oppenheimer for the data from the specially analyzed EAGLE halos shown in **Figure 9**. We also thank Lauren Corlies, Matt McQuinn, Andrew Fox, Romeel Davé, and John O'Meara for insightful comments on a draft of this article. We have made extensive use of NASA's Astrophysics Data System, *astropy* (Robitaille et al. 2013), *matplotlib* (Hunter 2007), *yt* (Turk et al. 2011), and the python tools *Colossus* from Benedikt Diemer and Seaborn by Michael Waskom.

LITERATURE CITED

- Adelberger KL, Steidel CC, Shapley AE, Pettini M. 2003. *Ap. J.* 584:45–75
Allen MG, Groves BA, Dopita MA, et al. 2008. *Ap. J. Suppl.* 178:20–55

- Anderson ME, Bregman JN. 2010. *Ap. J.* 714:320–31
- Anderson ME, Bregman JN. 2011. *Ap. J.* 737:22
- Anderson ME, Bregman JN, Dai X. 2013. *Ap. J.* 762:106
- Anderson ME, Churazov E, Bregman JN. 2016. *MNRAS* 455:227–43
- Armillotta L, Werk JK, Prochaska JX, et al. 2016. *MNRAS*. Submitted. arXiv:1608.05416
- Arrigoni Battaia F, Yang Y, Hennawi JF, et al. 2015. *Ap. J.* 804:26
- Bahcall JN, Spitzer L Jr. 1969. *Ap. J. Lett.* 156:L63
- Becker GD, Bolton JS, Lidz A. 2015. *Publ. Astron. Soc. Aust.* 32:e045
- Begelman MC, Fabian AC. 1990. *MNRAS* 244:26P–29P
- Behroozi PS, Conroy C, Wechsler RH. 2010. *Ap. J.* 717:379–403
- Benjamin RA. 1994. *The Origin and Evolution of Galactic Halo Gas*. Ph.D. thesis, Univ. Texas at Austin
- Bergeron J. 1986. *Astron. Astrophys.* 155:L8–11
- Bergeron J, Boissé P. 1991. *Astron. Astrophys.* 243:344–66
- Bergeron J, Stasińska G. 1986. *Astron. Astrophys.* 169:1–13
- Bertone S, Schaye J, Booth CM, et al. 2010. *MNRAS* 408:1120–38
- Binney J, Nipoti C, Fraternali F. 2009. *MNRAS* 397:1804–15
- Bland-Hawthorn J, Gerhard O. 2016. *Annu. Rev. Astron. Astrophys.* 54:529–96
- Bland-Hawthorn J, Sutherland R, Agertz O, Moore B. 2007. *Ap. J. Lett.* 670:L109–12
- Bogdán Á, Forman WR, Vogelsberger M, et al. 2013. *Ap. J.* 772:97
- Bordoloi R, Heckman TM, Norman CA. 2016. *Ap. J.* Submitted. arXiv:1605.07187
- Bordoloi R, Lilly SJ, Hardmeier E, et al. 2014a. *Ap. J.* 794:130
- Bordoloi R, Lilly SJ, Knobel C, et al. 2011. *Ap. J.* 743:10
- Bordoloi R, Tumlinson J, Werk JK, et al. 2014b. *Ap. J.* 796:136
- Borthakur S, Heckman T, Tumlinson J, et al. 2015. *Ap. J.* 813:46
- Bouché N, Murphy MT, Kacprzak GG, et al. 2013. *Science* 341:50–53
- Bowen DV, Chelouche D. 2011. *Ap. J.* 727:47
- Bowen DV, Chelouche D, Jenkins EB, et al. 2016. *Ap. J.* 826:50
- Brüns C, Kerp J, Kalberla PMW, Mebold U. 2000. *Astron. Astrophys.* 357:120–28
- Burchett JN, Tripp TM, Bordoloi R, et al. 2016. *Ap. J.* 832:124
- Cantalupo S, Arrigoni-Battaia F, Prochaska JX, et al. 2014. *Nature* 506:63–66
- Chen HW, Helsby JE, Gauthier JR, et al. 2010. *Ap. J.* 714:1521–41
- Chen HW, Lanzetta KM, Webb JK, Barcons X. 1998. *Ap. J.* 498:77–94
- Chen HW, Mulchaey JS. 2009. *Ap. J.* 701:1219–42
- Choi E, Ostriker JP, Naab T, et al. 2015. *MNRAS* 449:4105–16
- Christensen CR, Davé R, Governato F, et al. 2016. *Ap. J.* 824:57
- Churchill CW, Vander Vliet JR, Trujillo-Gomez S, et al. 2015. *Ap. J.* 802:10
- Cooksey KL, Kao MM, Simcoe RA, O’Meara JM, Prochaska JX. 2013. *Ap. J.* 763:37
- Cooksey KL, Thom C, Prochaska JX, Chen H-W. 2010. *Ap. J.* 708:868–908
- Cooper TJ, Simcoe RA, Cooksey KL, et al. 2015. *Ap. J.* 812:58
- Corlies L, Schiminovich D. 2016. *Ap. J.* 827:148
- Crighton NHM, O’Meara JM, Murphy MT. 2017. *MNRAS* 457:44
- Dai X, Anderson ME, Bregman JN, Miller JM. 2012. *Ap. J.* 755:107
- Dalcanton J, Seager S, Aigrain S, et al. 2015. AURA Rep., Assoc. Univ. Res. Astron., Washington, DC. arXiv:1507.04779
- Davé R, Finlator K, Oppenheimer BD. 2012. *MNRAS* 421:98–107
- Dekel A, Birnboim Y. 2006. *MNRAS* 368:2–20
- Dekel A, Mandelker N. 2014. *MNRAS* 444:2071–84
- Dekel A, Woo J. 2003. *MNRAS* 344:1131–44
- Di Matteo T, Colberg J, Springel V, et al. 2008. *Ap. J.* 676:33–53
- Dopita MA, Sutherland RS. 1996. *Ap. J. Suppl.* 102:161
- Edgar RJ, Chevalier RA. 1986. *Ap. J. Lett.* 310:L27–30
- Faerman Y, Sternberg A, McKee CF. 2017. *Ap. J.* 835:52
- Faucher-Giguere CA, Feldmann R, Quataert E, et al. 2016. *MNRAS* 461:L32–36

- Ferland GJ, Porter RL, van Hoof PAM, et al. 2013. *Rev. Mex. Astron. Astrofis.* 49:137–63
- Fielding D, Quataert E, McCourt M, Thompson TA. 2017. *MNRAS* 466:3810
- Finlator K, Davé R. 2008. *MNRAS* 385:2181–204
- Ford AB, Davé R, Oppenheimer BD, et al. 2014. *MNRAS* 444:1260–81
- Ford AB, Oppenheimer BD, Davé R, et al. 2013. *MNRAS* 432:89–112
- Ford AB, Werk JK, Davé R, et al. 2016. *MNRAS* 459:1745–63
- Fox A, Davé R, eds. 2017. *Gas Accretion onto Galaxies*, vol. 430. New York: Springer
- Fox AJ, Bordoloi R, Savage BD, et al. 2015. *Ap. J. Lett.* 799:L7
- Fox AJ, Lehner N, Tumlinson J, et al. 2013. *Ap. J.* 778:187
- Fox AJ, Prochaska JX, Ledoux C, et al. 2009. *Astron. Astrophys.* 503:731–46
- Fox AJ, Savage BD, Wakker BP. 2006. *Ap. J. Suppl.* 165:229–55
- Fox AJ, Wakker BP, Barger KA, et al. 2014. *Ap. J.* 787:147
- Fox AJ, Wakker BP, Savage BD, et al. 2005. *Ap. J.* 630:332–54
- Fraternali F, Binney JJ. 2008. *MNRAS* 386:935–44
- Fraternali F, Marasco A, Armillotta L, Marinacci F. 2015. *MNRAS* 447:L70–74
- Fumagalli M, Prochaska JX, Kasen D, et al. 2011. *MNRAS* 418:1796–821
- Gabor JM, Davé R. 2012. *MNRAS* 427:1816–29
- Gabor JM, Davé R, Finlator K, Oppenheimer BD. 2010. *MNRAS* 407:749–71
- Gallazzi A, Brinchmann J, Charlot S, White SDM. 2008. *MNRAS* 383:1439–58
- Gauthier JR, Chen HW, Tinker JL. 2010. *Ap. J.* 716:1263–68
- Geha M, Blanton MR, Yan R, Tinker JL. 2012. *Ap. J.* 757:85
- Glidden A, Cooper TJ, Cooksey KL, et al. 2016. *Ap. J.* 833:270
- Gnat O, Sternberg A. 2007. *Ap. J. Suppl.* 168:213–30
- Gnat O, Sternberg A. 2009. *Ap. J.* 693:1514–42
- Gnat O, Sternberg A, McKee CF. 2010. *Ap. J.* 718:1315–31
- Goerdt T, Dekel A, Sternberg A, et al. 2010. *MNRAS* 407:613–31
- Greco JP, Hill JC, Spergel DN, Battaglia N. 2015. *Ap. J.* 808:151
- Grimes JP, Heckman T, Aloisi A, et al. 2009. *Ap. J. Suppl.* 181:272–320
- Gupta A, Mathur S, Krongold Y, et al. 2012. *Ap. J. Lett.* 756:L8
- Gutcke TA, Stinson GS, Macciò AV, et al. 2017. *MNRAS* 464:2796–815
- Haardt F, Madau P. 2001. In *Clusters of Galaxies and the High Redshift Universe Observed in X-rays. Recent results of XMM-Newton and Chandra, XXXVth Rencontres de Moriond, XXIst Moriond Astrophys. Meet., March 10–17, Savoie, France*, ed. DM Neumann, F Durret, JTT Van. arXiv:astro-ph/0106018
- Hafen Z, Faucher-Giguere CA, Angles-Alcázar D, et al. 2016. *MNRAS*. Submitted. arXiv:1608.05712
- Hayes M, Melinder J, Östlin G, et al. 2016. *Ap. J.* 828:49
- Heckman TM, Alexandroff RM, Borthakur S, et al. 2015. *Ap. J.* 809:147
- Heckman TM, Borthakur S. 2016. *Ap. J.* 822:9
- Heckman TM, Norman CA, Strickland DK, Sembach KR. 2002. *Ap. J.* 577:691–700
- Heitsch F, Putman ME. 2009. *Ap. J.* 698:1485–96
- Henry A, Scarlata C, Martin CL, Erb D. 2015. *Ap. J.* 809:19
- Hill JC, Ferraro S, Battaglia N, et al. 2016. *Phys. Rev. Lett.* 117:1301
- Hopkins PF, Hernquist L, Cox TJ, et al. 2006. *Ap. J.* 639:700–9
- Hopkins PF, Quataert E, Murray N. 2012. *MNRAS* 421:3522–37
- Hummels C, Smith B, Silvia D. 2016. *Ap. J.* Submitted. arXiv:1612.03935
- Hummels CB, Bryan GL, Smith BD, Turk MJ. 2013. *MNRAS* 430:1548–65
- Humphrey PJ, Buote DA, Canizares CR, et al. 2011. *Ap. J.* 729:53
- Hunter JD. 2007. *Comput. Sci. Eng.* 9:90–95
- Johnson SD, Chen HW, Mulchaey JS. 2015a. *MNRAS* 449:3263–73
- Johnson SD, Chen HW, Mulchaey JS. 2015b. *MNRAS* 452:2553–65
- Johnson SD, Chen HW, Mulchaey JS, et al. 2014. *MNRAS* 438:3039–48
- Kacprzak GG, Churchill CW, Steidel CC, et al. 2012. *MNRAS* 427:3029–43
- Keel WC, Lintott CJ, Schawinski K, et al. 2012. *Astron. J.* 144:66
- Kereš D, Hernquist L. 2009. *Ap. J. Lett.* 700:L1–5

- Kereš D, Katz N, Weinberg DH, Davé R. 2005. *MNRAS* 363:2–28
- Kollmeier JA, Weinberg DH, Oppenheimer BD, et al. 2014. *Ap. J. Lett.* 789:L32
- Kornei KA, Shapley AE, Martin CL, et al. 2012. *Ap. J.* 758:135
- Kwak K, Shelton RL. 2010. *Ap. J.* 719:523–39
- Lanzetta KM, Bowen DV, Tytler D, Webb JK. 1995. *Ap. J.* 442:538–68
- Lehner N, Howk JC. 2011. *Science* 334:955
- Lehner N, Howk JC, Tripp TM, et al. 2013. *Ap. J.* 770:138
- Lehner N, Howk JC, Wakker BP. 2015. *Ap. J.* 804:79
- Lehner N, O’Meara JM, Fox AJ, et al. 2014. *Ap. J.* 788:119
- Lehner N, O’Meara JM, Howk JC, et al. 2016. *Ap. J.* 833:283
- Lehner N, Prochaska JX, Kobulnicky HA, et al. 2009. *Ap. J.* 694:734–50
- Lehnert MD, Le Tiran L, Nesvadba NPH, et al. 2013. *Astron. Astrophys.* 555:A72
- Liang CJ, Chen HW. 2014. *MNRAS* 445:2061–81
- Lilly SJ, Carollo CM, Pipino A, et al. 2013. *Ap. J.* 772:119
- Lu Y, Blanc GA, Benson A. 2015. *Ap. J.* 808:129
- Maller AH, Bullock JS. 2004. *MNRAS* 355:694–712
- Martin CL. 2005. *Ap. J.* 621:227–45
- Martin CL, Shapley AE, Coil AL, et al. 2012. *Ap. J.* 760:127
- Martin DC, Matuszewski M, Morrissey P, et al. 2015. *Nature* 524:192–95
- Matejek MS, Simcoe RA. 2012. *Ap. J.* 761:112
- Mathes NL, Churchill CW, Kacprzak GG, et al. 2014. *Ap. J.* 792:128
- McCourt M, Sharma P, Quataert E, Parrish IJ. 2012. *MNRAS* 419:3319–37
- McGaugh SS, Schombert JM, de Blok WJG, Zagursky MJ. 2010. *Ap. J. Lett.* 708:L14–17
- Meiring JD, Tripp TM, Werk JK, et al. 2013. *Ap. J.* 767:49
- Ménard B, Scranton R, Fukugita M, Richards G. 2010. *MNRAS* 405:1025–39
- Mo HJ, Miralda-Escudé J. 1996. *Ap. J.* 469:589
- Morton DC. 2003. *Ap. J. Suppl.* 149:205–38
- Münch G, Zirin H. 1961. *Ap. J.* 133:11
- Muratov AL, Keres D, Faucher-Giguère CA, et al. 2015. *MNRAS* 454:2691–713
- Muratov AL, Keres D, Faucher-Giguère CA, et al. 2017. *MNRAS* 468:4170–88
- Murray N, Ménard B, Thompson TA. 2011. *Ap. J.* 735:66
- Murray N, Quataert E, Thompson TA. 2005. *Ap. J.* 618:569–85
- Muzahid S, Kacprzak GG, Churchill CW, et al. 2015. *Ap. J.* 811:132
- Muzahid S, Srianand R, Bergeron J, Petitjean P. 2012. *MNRAS* 421:446–67
- Narayanan A, Savage BD, Wakker BP, et al. 2011. *Ap. J.* 730:15
- Narayanan A, Wakker BP, Savage BD, et al. 2010. *Ap. J.* 721:960–74
- Nelson D, Vogelsberger M, Genel S, et al. 2013. *MNRAS* 429:3353–70
- Nicastro F, Mathur S, Elvis M, et al. 2005. *Nature* 433:495–98
- Nielsen NM, Churchill CW, Kacprzak GG, Murphy MT. 2013. *Ap. J.* 776:114
- Nielsen NM, Churchill CW, Kacprzak GG, et al. 2015. *Ap. J.* 812:83
- Nielsen NM, Churchill CW, Kacprzak GG, et al. 2016. *Ap. J.* 818:171
- Oppenheimer BD, Crain RA, Schaye J, et al. 2016. *MNRAS* 460:2157–79
- Oppenheimer BD, Davé R. 2006. *MNRAS* 373:1265–92
- Oppenheimer BD, Davé R. 2008. *MNRAS* 387:577–600
- Oppenheimer BD, Davé R, Kereš D, et al. 2010. *MNRAS* 406:2325–38
- Oppenheimer BD, Schaye J. 2013a. *MNRAS* 434:1043–62
- Oppenheimer BD, Schaye J. 2013b. *MNRAS* 434:1063–78
- Peek JEG, Ménard B, Corrales L. 2015. *Ap. J.* 813:7
- Peebles MS, Werk JK, Tumlinson J, et al. 2014. *Ap. J.* 786:54
- Planck Collab., Ade PAR, Aghanim N, et al. 2013. *Astron. Astrophys.* 557:A52
- Prescott MKM, Martin CL, Dey A. 2015. *Ap. J.* 799:62
- Prochaska J, Werk JK, Worsek G, et al. 2017. *Ap. J.* 837:169
- Prochaska JX, Bloom JS, Chen HW, et al. 2004. *Ap. J.* 611:200–7

- Prochaska JX, Lau MW, Hennawi JF. 2014. *Ap. J.* 796:140
- Prochaska JX, Weiner B, Chen HW, et al. 2011a. *Ap. J.* 740:91
- Prochaska JX, Weiner B, Chen HW, et al. 2011b. *Ap. J. Suppl.* 193:28
- Putman ME, Peek JEG, Jounge MR. 2012. *Annu. Rev. Astron. Astrophys.* 50:491–529
- Rauch M, Haehnelt MG. 2011. *MNRAS* 412:L55–57
- Rauch M, Sargent WLW, Barlow TA. 2001. *Ap. J.* 554:823–40
- Richter P, Savage BD, Tripp TM, Sembach KR. 2004. *Ap. J. Suppl.* 153:165–204
- Robitaille TP, Tollerud EJ, Greenfield P, et al. 2013. *Astron. Astrophys.* 558:A33
- Rubin KHR, Hennawi JF, Prochaska JX, et al. 2015. *Ap. J.* 808:38
- Rubin KHR, Prochaska JX, Koo DC, Phillips AC. 2012. *Ap. J. Lett.* 747:L26
- Rubin KHR, Prochaska JX, Koo DC, et al. 2014. *Ap. J.* 794:156
- Rudie GC, Steidel CC, Trainor RF, et al. 2012. *Ap. J.* 750:67
- Salem M, Bryan GL, Corlies L. 2016. *MNRAS* 456:582–601
- Sargent WLW, Young PJ, Boksenberg A, Tytler D. 1980. *Ap. J. Suppl.* 42:41–81
- Savage BD, Kim TS, Wakker BP, et al. 2014. *Ap. J. Suppl.* 212:8
- Savage BD, Lehner N, Narayanan A. 2011. *Ap. J.* 743:180
- Schawinski K, Urry CM, Simmons BD, et al. 2014. *MNRAS* 440:889–907
- Schaye J, Carswell RF, Kim TS. 2007. *MNRAS* 379:1169–94
- Schaye J, Crain RA, Bower RG, et al. 2015. *MNRAS* 446:521–54
- Schiminovich D, Catinella B, Kauffmann G, et al. 2010. *MNRAS* 408:919–34
- Sembach KR, Tripp TM, Savage BD, Richter P. 2004. *Ap. J. Suppl.* 155:351–93
- Sembach KR, Wakker BP, Savage BD, et al. 2003. *Ap. J. Suppl.* 146:165–208
- Shapiro PR, Field GB. 1976. *Ap. J.* 205:762–65
- Shattow GM, Croton DJ, Bibiano A. 2015. *MNRAS* 450:2306–16
- Shen S, Madau P, Aguirre A, et al. 2012. *Ap. J.* 760:50
- Shen S, Madau P, Guedes J, et al. 2013. *Ap. J.* 765:89
- Silvia DW. 2013. PhD Thesis, University of Colorado, Boulder
- Slavin JD, Shull JM, Begelman MC. 1993. *Ap. J.* 407:83–99
- Somerville RS, Davé R. 2015. *Annu. Rev. Astron. Astrophys.* 53:51–113
- Somerville RS, Popping G, Trager SC. 2015. *MNRAS* 453:4337–67
- Spitzer L Jr. 1956. *Ap. J.* 124:20
- Springel V, Di Matteo T, Hernquist L. 2005. *MNRAS* 361:776–94
- Steidel CC, Erb DK, Shapley AE, et al. 2010. *Ap. J.* 717:289–322
- Stern J, Hennawi JF, Prochaska JX, Werk JK. 2016. *Ap. J.* 830:87
- Stewart KR, Kaufmann T, Bullock JS, et al. 2011. *Ap. J.* 738:39
- Stinson GS, Brook C, Prochaska JX, et al. 2012. *MNRAS* 425:1270–77
- Stocke JT, Keeney BA, Danforth CW, et al. 2013. *Ap. J.* 763:148
- Stocke JT, Penton SV, Danforth CW, et al. 2006. *Ap. J.* 641:217–28
- Strickland DK, Heckman TM. 2009. *Ap. J.* 697:2030–56
- Suresh J, Rubin KHR, Kannan R, et al. 2017. *MNRAS* 465:2966–82
- Tejos N, Prochaska JX, Crighton NHM, et al. 2016. *MNRAS* 455:2662–97
- Tepper-García T, Bland-Hawthorn J, Sutherland RS. 2015. *Ap. J.* 813:94
- Thom C, Tumlinson J, Werk JK, et al. 2012. *Ap. J. Lett.* 758:L41
- Thompson TA, Quataert E, Zhang D, Weinberg DH. 2016. *MNRAS* 455:1830–44
- Tinsley BM. 1980. *Fundam. Cosm. Phys.* 5:287–388
- Tremonti CA, Heckman TM, Kauffmann G, et al. 2004. *Ap. J.* 613:898–913
- Tripp TM, Giroux ML, Stocke JT, et al. 2001. *Ap. J.* 563:724–35
- Tripp TM, Meiring JD, Prochaska JX, et al. 2011. *Science* 334:952
- Tripp TM, Sembach KR, Bowen DV, et al. 2008. *Ap. J. Suppl.* 177:39–102
- Tumlinson J, Shull JM, Giroux ML, Stocke JT. 2005. *Ap. J.* 620:95–112
- Tumlinson J, Thom C, Werk JK, et al. 2011. *Science* 334:948
- Tumlinson J, Thom C, Werk JK, et al. 2013. *Ap. J.* 777:59
- Turk MJ, Smith BD, Oishi JS, et al. 2011. *Ap. J. Suppl.* 192:9

- Turner ML, Schaye J, Crain RA, et al. 2016. *MNRAS* 462:2440–64
- Turner ML, Schaye J, Steidel CC, et al. 2014. *MNRAS* 445:794–822
- Turner ML, Schaye J, Steidel CC, et al. 2015. *MNRAS* 450:2067–82
- van de Voort F, Schaye J, Altay G, Theuns T. 2012. *MNRAS* 421:2809–19
- Veilleux S, Cecil G, Bland-Hawthorn J. 2005. *Annu. Rev. Astron. Astrophys.* 43:769–826
- Vikhlinin A, Kravtsov A, Forman W, et al. 2006. *Ap. J.* 640:691–709
- Vogelsberger M, Genel S, Springel V, et al. 2014. *MNRAS* 444:1518–47
- Voit GM, Bryan GL, O’Shea BW, Donahue M. 2015a. *Ap. J. Lett.* 808:L30
- Voit GM, Donahue M, Bryan GL, McDonald M. 2015b. *Nature* 519:203–6
- Wakker BP, Hernandez AK, French DM, et al. 2015. *Ap. J.* 814:40
- Wakker BP, Savage BD, Fox AJ, et al. 2012. *Ap. J.* 749:157
- Walker SA, Bagchi J, Fabian AC. 2015. *MNRAS* 449:3527–34
- Wang B. 1995. *Ap. J.* 444:590–609
- Wang QD, Yao Y. 2012. arXiv:1211.4834
- Werk JK, Prochaska JX, Cantalupo S, et al. 2016. *Ap. J.* 833:54
- Werk JK, Prochaska JX, Thom C, et al. 2012. *Ap. J. Suppl.* 198:3
- Werk JK, Prochaska JX, Thom C, et al. 2013. *Ap. J. Suppl.* 204:17
- Werk JK, Prochaska JX, Tumlinson J, et al. 2014. *Ap. J.* 792:8
- Whitaker KE, van Dokkum PG, Brammer G, Franx M. 2012. *Ap. J. Lett.* 754:L29
- Williams RJ, Mathur S, Nicastro F, et al. 2005. *Ap. J.* 631:856–67
- Wotta CB, Lehner N, Howk JC, et al. 2016. *Ap. J.* 831:95
- Wu KKS, Fabian AC, Nulsen PEJ. 2001. *MNRAS* 324:95–107
- Yao Y, Shull JM, Wang QD, Cash W. 2012. *Ap. J.* 746:166
- Yao Y, Wang QD, Penton SV, et al. 2010. *Ap. J.* 716:1514–21
- York DG, Khare P, Vanden Berk D, et al. 2006. *MNRAS* 367:945–78
- Zahid HJ, Bresolin F, Kewley LJ, et al. 2012. *Ap. J.* 750:120
- Zhang H, Zaritsky D, Zhu G, et al. 2016. *Ap. J.* 833:276
- Zheng Y, Peek JEG, Werk JK, Putman ME. 2017. *Ap. J.* 834:179
- Zhu G, Ménard B. 2013a. *Ap. J.* 770:130
- Zhu G, Ménard B. 2013b. *Ap. J.* 773:16
- Zhu G, Ménard B, Bizyaev D, et al. 2014. *MNRAS* 439:3139–55
- Zhu W, Feng L, Xia Y, et al. 2013. *Ap. J.* 777:48
- Zu Y, Weinberg DH, Davé R, et al. 2011. *MNRAS* 412:1059–69